

OBA Working Paper Series

Paper No. 5, June 2005

Techniques for estimating the fiscal costs and risks of long-term output-based payments

Glenn Boyle and Timothy Irwin*

June 2005

Long-term commitments to make output-based payments for infrastructure can encourage private investors to provide socially valuable services. Making good decisions about such commitments is difficult, however, unless the government understands the fiscal costs and risks of possible commitments. Considering voucher schemes, shadow tolls, availability payments, and access, connection, and consumption subsidies, this paper considers measures of the fiscal risks of such commitments, including the excess-payment probability and cash-flow-at-risk. Then it illustrates

techniques, based on modern finance theory, for valuing payment commitments by taking account of the timing of payments and their risk-characteristics. Although the paper is inevitably mathematical, it focuses on practical applications and shows how the techniques can be implemented in spreadsheets.

* **Glenn Boyle** is the Executive Director of the New Zealand Institute for the Study of Competition and Regulation at Victoria University of Wellington (glenn.boyle@vuw.ac.nz).
Timothy Irwin is in the Infrastructure Advisory Services group at the World Bank (tirwin@worldbank.org).

About the Global Partnership on Output-Based Aid

What is output-based aid? OBA is a strategy for supporting the delivery of basic services—water, sanitation, electricity, telecommunications, transport, education, health care—where policy concerns would justify public funding to complement or replace user fees.

How can GPOBA help? GPOBA can assist in the design and development of pilot OBA projects, and can help identify and disseminate emerging knowledge on issues related to OBA approaches through studies, publications, workshops, and conferences. GPOBA can also contribute to the funding of subsidized payments for the provision of services under OBA schemes.

Global Partnership on Output Based Aid

World Bank
Mailstop: H3-300
Washington, DC 20433, USA

To find out more, visit
www.gpoba.org

Techniques for estimating the fiscal costs and risks of long-term output-based payments

Glenn Boyle and Timothy Irwin*

June 2005

* Glenn Boyle is the Executive Director of the New Zealand Institute for the Study of Competition and Regulation at Victoria University of Wellington (glenn.boyle@vuw.ac.nz). Timothy Irwin is in the Infrastructure Advisory Services group at the World Bank (tirwin@worldbank.org).

Abstract

Long-term commitments to make output-based payments for infrastructure can encourage private investors to provide socially valuable services. Making good decisions about such commitments is difficult, however, unless the government understands the fiscal costs and risks of possible commitments. Considering voucher schemes, shadow tolls, availability payments, and access, connection, and consumption subsidies, this paper considers measures of the fiscal risks of such commitments, including the excess-payment probability and cash-flow-at-risk. Then it illustrates techniques, based on modern finance theory, for valuing payment commitments by taking account of the timing of payments and their risk-characteristics. Although the paper is inevitably mathematical, it focuses on practical applications and shows how the techniques can be implemented in spreadsheets.

Contents

1	Introduction	9
2	Output-based payments	11
3	Risk measurement	15
3.1	A simple model of risky variables.....	15
3.2	Geometric brownian motion—a useful special case	16
3.3	The standard deviation as a measure of risk	18
3.4	Excess payment probability	19
3.5	Cash-flow-at-risk	26
3.6	Estimating the entire probability distribution.....	29
4	General principles for valuing output-based schemes	32
4.1	Adjusting the discount rate	33
4.2	Estimating certainty equivalents	37
4.3	Equivalence in principle of the two methods	38
4.4	Superiority in practice of the certainty-equivalent method	40
5	Valuation applications	43
5.1	Uncapped annual-output schemes	43
5.2	Capped annual-output schemes.....	45
5.3	Schemes with multiple payment rates.....	47
5.4	Total cumulative cap	51
5.5	Access subsidies and multi-year voucher schemes	53
6	Some additional valuation issues	55
6.1	Varying growth rates	55
6.2	Mean-reverting processes.....	55
6.3	Non-CAPM risk premiums	56
	References	63

Examples

Example 1: Uncapped utility connection subsidy—estimating the excess payment probability (annual expenditure) using an analytical approach	21
Example 2: Uncapped utility connection subsidy—estimating the excess payment probability (total expenditure) using Monte Carlo simulation	24
Example 3: Uncapped utility-connection subsidy—estimating cash-flow-at-risk using Monte Carlo simulation	27
Example 4: Uncapped utility connection subsidy—valuing the subsidy using an analytical approach	43
Example 5: Annually capped utility connections subsidy—valuing the subsidy using an analytical approach	47
Example 6: Scheme with multiple payment rates—valuation using Monte Carlo simulation	49
Example 7: Utility connections subsidy with cumulative cap—valuing the subsidy using Monte Carlo simulation	52

1 Introduction

Output-based payments are an important tool of infrastructure policy. Sometimes, governments subsidize services sold to households. At other times, governments are the sole source of a private infrastructure firm's revenue. When the government is the sole purchaser, the arrangements are sometimes called public-private partnerships. When the government subsidizes services sold to end users, the subsidy is sometimes called output-based aid.¹ In all cases, the government pays only when the firm delivers a service (when a connection is made, a car uses a road, or power is made available).

Various challenges arise in the design of output-based payments, such as defining outputs, monitoring performance, and making credible the government's commitment to make payments in the future. Another problem is to estimate the fiscal cost of commitments to pay *a priori* unknown amounts and the fiscal risks the government faces because of the payments—and to compare these costs and risks with the costs and risks of alternative forms of subsidy. In this report, we outline techniques governments can use to make these estimates, considering a wide range of output-based payments.

When the government commits itself to payments for only a year, the fiscal risks are likely to be small. But if the payments are to encourage service providers to invest in order to provide services,

¹ For more information on output-based aid in general, see Brook and Petrie (2001); Smith (2001); and information on the website of the Global Partnership for Output Aid www.gpoba.org. For discussion of its use in water, see Drees, Schwartz, and Bakalion (2004); Gómez-Lobo (2001); and Marin (2002). For electricity, see Harris (2002) and Tomkins (2001). For transport, see Liautaud (2001) and Scott and Birnie (2002). For telecommunications, see Cannock (2002). Many of these papers are collected in Brook and Smith (2001) and on the GPOBA website.

the government may have to commit itself in advance to offering the payments for many years—perhaps for as long as the life of the assets used to provide the service. Even in this case, if the amounts of money are small or not subject to much risk, there may not be a strong case for carefully measuring the fiscal risks the government is taking and valuing the obligations it is incurring. When the payments represent long-term commitments of potentially large and uncertain amounts, however, the techniques set out here may help the government understand the costs and risks associated with the decisions it is making.

After briefly describing various types of output-based payment scheme, the report discusses how governments can *measure* the fiscal risks associated with such payments. How can financial models incorporate random variation? And how should *risk* be defined? The next sections then consider how the government can *value* its exposure to risk. That is, given an estimate of the risks surrounding future payments, how can the government estimate the present value of its obligation to make the payments? We start this discussion by describing general principles for valuing payment commitments, then apply those principles to different types of output-based payment, and finally consider some more-complex issues that may arise.

The report is inevitably mathematical: it isn't possible to say much about how a government can measure or value its exposure to risk without talking about, for example, the formulas that need to be entered into spreadsheets. Yet the approach we follow emphasizes economic intuition rather than mathematical purity.² We also include seven worked examples intended to clarify how the

² For readers who are interested in more explanation or a more rigorous treatment, we provide some suggestions for appropriate reading.

approach can be implemented. An accompanying spreadsheet is available on request.³

2 Output-based payments

Output-based payments come in many varieties. We mention just a few below.

Consumption subsidies. Governments can subsidize the consumption of services such as water or electricity. In Chile, for example, the government subsidizes the first five cubic meters a month of water consumption by eligible customers (Gómez-Lobo, 2002). Uncertainty about cost of consumption subsidies can arise from uncertainty about consumption per subsidized customer and the number of customers eligible for subsidies.

Voucher schemes. In voucher schemes, the government makes a payment to a customer to help the customer pay for social services, such as education or health. Education vouchers, for example, may pay a fixed amount per year toward the cost of education for, say, three years. Uncertainty about the cost can arise because of uncertainty about the number of students eligible for subsidies and their propensity to enroll in the subsidized programs.

Connection subsidies. In a connection-subsidy scheme, the government agrees to pay a utility for each new eligible customer it connects to the network.⁴ In Guatemala, for example, the government used such a scheme to increase rural electrification when it privatized the two companies serving rural areas. It agreed to pay US\$650 for each verified connection made to an eligible

³ Email Tim Irwin at tirwin@worldbank.org

⁴ Connection may seem more like an input used to provide the output of access to services than an output itself. In practice, the term output-based subsidies is used quite broadly to include connection subsidies, in part reflecting the very close connection between connection and access.

household, and between May 1999 and May 2002, 122,000 subsidized connections were made (Harris, 2002). In the water sector in Paraguay, the government has piloted an output-based aid scheme that pays US\$150 per new connection up to a maximum number of connections (Drees, Schwartz, and Bakalian, 2004). Connection subsidies are also possible in natural-gas distribution and telecommunications. Uncertainty about cost can arise from uncertainty about the demand for and the supply of new connections, as well as about the number of households meeting eligibility criteria.

Access subsidies. Access subsidies are like connection subsidies, except that the government pays a subsidy for each year in which the connection to the eligible customer is in place, not just for the creation of the connection. Uncertainty about their cost can arise from the same sources of uncertainty that affect connection subsidies, as well as uncertainty about whether connections will be maintained over time.

Availability payments. Governments sometimes agree to purchase capacity to produce an output, such as electricity or drinkable water. For example, they may agree (typically through a state-owned utility) to make a series of availability payments to an independent power company if the power company has the capacity to produce power. Similar contracts are being used to procure education and health facilities in the United Kingdom and elsewhere under the name of the “private finance initiative” and “public-private partnerships.”

Shadow-tolls. Under a shadow toll, the government agrees to make traffic-dependent payments to a road company. The payments might be a constant amount per vehicle, or they might be a declining stepwise function of the number of vehicles using the road. For example, one shadow-toll road in the UK (see National Audit Office 1998) has four bands; tolls per vehicle decline as traffic increases and in the fourth band are zero (so there is a maximum

payment). Uncertainty about the cost in such a case depends largely on uncertain about future traffic flows.

Table 1 summarizes.

Table 1: Types of output-based payment scheme

Type	Possible usage	Risk
Consumption subsidies	Water, electricity	Consumption per subsidized customer, number of eligible customers
Vouchers	Education, health	Number of eligible customers, propensity to enroll
Connection subsidies	Water, electricity, gas, telecommunications	Demand for new connections, supply of new connections, number of eligible customers
Access subsidies	Water, electricity, gas, telecommunications	Propensity of customers to maintain access (as well as factors for connection subsidies)
Availability payments	Wholesale water and electricity, roads, and school, hospital, and prison facilities	Supply of capacity
Shadow tolls	Roads	Traffic flows

Output-based payments are many and various, but for the purpose of measuring and valuing the government's exposure to risk, not all the differences are crucial. What matters most for analytical purposes is whether the payments depend on current output or cumulative output and whether expenditure is capped or uncapped.

Many of the payment schemes have the same payment structure

$$Y_t = sX_t,$$

where X is the output from some economic activity and s is the dollar amount paid for each unit of output. Examples of this sort of scheme include connection subsidies, consumption subsidies, and availability-payment schemes:

With a connection subsidy, X is the number of new connections.

With a consumption scheme, X is the volume of consumption.

With an availability-payment scheme, X is the capacity made available.

In each case, the payment in year t depends only on the level of X during that year, so we group these under the general heading of annual-output schemes. (We use the yearly structure purely for expositional convenience. Exactly the same principles apply if payments are made on some other periodic basis, such as every month or every quarter.)

In other cases, the schemes can best be analyzed as subsidizing some form of *cumulative* output. Access might be analyzed as the cumulative result of previous connections (at least when the probability of disconnection is low). In this case the subsidy depends not just on this year's rate of connections, but also on past years' rates of connection. Voucher schemes applying to a multi-year program might be analyzed similarly.

A further distinction arises between schemes on which expenditure is capped and those on which it is open ended. Under a capped scheme, the government places a ceiling on the number of outputs it will pay for or subsidize. Under an open-ended, or uncapped, scheme, it imposes no such limit. Capped schemes limit the government's costs and (downside) risk, but complicate the analysis.

3 Risk measurement

A first step in understanding the costs and risks of output-based payments is to estimate the likely amounts and variability of future expenditure. To do that, we need a way of modeling the economic variables determining the payments that allows for the variables to be risky: to fluctuate randomly as well as having a trend.

3.1 A simple model of risky variables

In its most general sense, stating that the level of some activity is risky simply means that future realizations of this activity are unknown, and therefore uncertain. For example, while the activity may have a trend, it can also have a random element, in which case its evolution is characterized by unpredictable fluctuations around the trend. One useful way of capturing this idea in a formal sense is to assume that the activity evolves as an *Itô process*. (For further reading, at a fairly intuitive level, on Itô processes, see Ritchken 1987, chapter 16 or Dixit and Pindyck 1994, chapter 3.)

To see what this means, suppose that W_t is the level or value of some variable W at date t ; for example, W could be the number of new phone connections made in a year under an output-based aid scheme, or it could be the dollar value of the payments made as a result of these new connections. If W follows an Itô process, then changes in W are given by

$$dW = a(W,t)dt + b(W,t)dZ \quad (1)$$

where $a(W,t)$ and $b(W,t)$ are deterministic (that is, known) functions of W and time t , and Z is a normally distributed (that is, “bell-shaped”) random variable with zero mean and variance equal to one. The last term in (1), $dZ \equiv Z\sqrt{dt}$, is a random variable (known formally as the increment of a Wiener process) that is normally distributed with zero mean and variance dt . In simple

terms, equation (1) states that changes in W at each instant in time have a trend (*drift*) component— $a(W,t)dt$ —and a stochastic (*diffusion*) component— $b(W,t)dZ$.

To those unfamiliar with the modeling of risk, these new terms and equations may seem daunting. Yet the discussion below, including the worked examples, shows how the ideas can be implemented in a relatively straightforward way in a spreadsheet.

3.2 Geometric brownian motion—a useful special case

A particularly convenient special case of (1) occurs when $a(W,t) = \mu W$ and $b(W,t) = \sigma W$, where μ and σ are constants. Then (1) becomes

$$\frac{dW}{W} = \mu dt + \sigma dZ \quad (2)$$

which states that percentage changes in W are normally distributed (because the random component dZ has a normal distribution) with an expected value of μ (per unit of time) and a standard deviation of σ (also per unit of time).⁵ In this case, W is said to follow a *geometric brownian motion*. Intuitively, expected growth in W over a unit of time ($dt = 1$) is expected to be constant at μ , but this is subject to random deviations that are proportional to σ .

Apart from being relatively straightforward to work with, equation (2) has the desirable property that W cannot be negative, a feature

⁵ This characterization of (2) is designed to capture its intuitive spirit rather than provide any rigorous description. For more on the latter, see Duffie (1996, ch 5).

common to many economic activities. In addition, if changes in W satisfy (2), then changes in the natural logarithm of W are given by⁶

$$d \ln W = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dZ, \quad (2a)$$

and, therefore, for $dt = 1$

$$\ln W_t - \ln W_{t-1} = \left(\mu - \frac{\sigma^2}{2} \right) + \sigma Z_t.$$

By adding $\ln W_{t-1}$ to both sides and then exponentiating, this can be written as

$$W_t = W_{t-1} \exp \left\{ \left(\mu - \frac{\sigma^2}{2} \right) + \sigma Z_t \right\}. \quad (3)$$

Equation (3) gives the value of W as a function of the parameter values μ and σ , and the unknown realization of Z .

Information about μ and σ is, obviously, of paramount importance to any attempted quantification of scheme risk and, as we shall see, cost. The best procedure for estimating these parameters will vary from case to case, and advice may well be required from experts in the area, such as statisticians and economic forecasters. Nevertheless, an approach that works in some cases is to assume that the future will look much like the past and, accordingly,

⁶ This follows from a mathematical result known as Itô's Lemma. For a discussion, see, for example, Hull (2003). Briefly, Itô's Lemma states that if W follows an Itô process, as set out in (1), changes in a function V of W are given by

$$dV = \left(\frac{\partial V}{\partial t} + a(W,t) \frac{\partial V}{\partial W} + \frac{b(W,t)^2}{2} \frac{\partial^2 V}{\partial W^2} \right) dt + b(W,t) \frac{\partial V}{\partial W} dZ.$$

Equation (2a) follows from the assumption of geometric Brownian motion, which implies that $a(W,t) = \mu W$ and $b(W,t) = \sigma W$, and by setting $V = \ln W$, so that $\partial V / \partial W = 1/W$, $\partial^2 V / \partial W^2 = -1/W^2$, and $\partial V / \partial t = 0$.

calculate the historical values of the parameters from available data. For example, given some data on past growth rates in W , μ can be estimated using the arithmetical average (the sample mean) of these data. When historical data are not available, financial analyses including forecast rates of growth may be available. Forecasts of volatility are less likely to be available and in such cases the best approach is often to look for historical data on similar projects.

3.3 The standard deviation as a measure of risk

Equations (1)–(3) capture the idea that future values of W are uncertain, and thus allow us to infer something about the risk of W . However, to do so, we first need to decide what aspects of the uncertainty inherent in W constitute risk; that is, we need a *definition* of risk.

The standard approach in finance theory emphasizes the role of σ , the standard deviation of growth in W . According to this view, the greater is σ , the greater is the risk of W . However, knowing that the standard deviation of growth in payments is, say, 10 percent, is not in itself very illuminating. Moreover, such a ranking leads to some counter-intuitive outcomes. Suppose there are two subsidy schemes for which expenditure evolves according to (3) with μ and σ as follows

$$\text{Scheme A: } \mu_A = 0.25; \quad \sigma_A = 0.08$$

$$\text{Scheme B: } \mu_B = 0.05; \quad \sigma_B = 0.12$$

Thus, scheme B has a higher σ than scheme A and might therefore be said to be riskier. But a “bad” outcome for scheme B (where actual growth in expenditure exceeds expected growth by one standard deviation) results in growth of 17 percent, which is exactly the same as that generated by a “good” outcome for scheme A (where actual

growth falls short of expected growth by one standard deviation). In this example, scheme A has lower σ than scheme B, but nevertheless exposes the government to a higher payment in most states. (We use “good” and “bad” here to refer to fiscal outcomes; obviously, to the extent that a subsidy is designed to encourage some activity, high levels of that activity need not be bad in a fundamental sense.)

The reason σ does not work very well as a measure of risk by itself in this case (and more generally) is because it is essentially a measure of volatility whereas risk, as Olsen (1997) points out, is usually thought of as the loss of something that one values, or so-called “downside risk”. In the remainder of this section, we shall explain and illustrate two measures of downside risk that are most relevant to the providers of an output-based payment.

3.4 Excess payment probability

Governments may be relatively indifferent about the size of potential payments so long as these do not exceed a certain level. This might be the case if, for example, the government’s fiscal position is threatened only by particularly high payments. One potentially useful measure of a scheme’s risk in these circumstances is the probability of payments exceeding some specified level. To illustrate this, let Y_t denote the date t cash payment made by the government to a private investor in return for the provision of some agreed service. For example, if the government offers a connection subsidy, Y_t equals the number of new connections at date t multiplied by the subsidy per unit. Then if Y_{up} is an upper bound on payments that the government is comfortable with making, the excess payment probability (*EPP*) is given by

$$EPP = \Pr(Y \geq Y_{up})$$

The excess payment probability tells us the likelihood of subsidy payments exceeding some upper threshold. All else equal, a subsidy scheme with a high excess payment probability is riskier than a scheme with a low excess payment probability.

3.4.1 Lognormally distributed payments

In some situations, it is possible to calculate a scheme's excess payment probability in a simple manner. In particular, if Y has a normal (bell-shaped) distribution, then knowledge of the mean and standard deviation of Y is sufficient for a spreadsheet program such as Excel to provide a quick calculation of the excess payment probability. The main problem with this approach is that Y will seldom have a normal distribution, as such a distribution allows for negative values. Instead, it is usually more reasonable to assume that Y follows a geometric brownian process (equation (2)), which permits only positive values.

Fortunately, if Y is indeed given by (2), then essentially the same procedure can be applied. Note first that $Y \geq Y_{up}$ if and only if $\ln Y \geq \ln Y_{up}$. Hence

$$\begin{aligned} EPP_t &= \Pr(Y_t \geq Y_{up}) \\ &= \Pr(\ln Y_t \geq \ln Y_{up}) \\ &= 1 - \Pr(\ln Y_t \leq \ln Y_{up}) \end{aligned}$$

Now note (from (3)) that if Y follows a geometric brownian motion, $\ln Y$ is normally distributed. Any normally distributed variable is a simple linear function of the standard normal variable Z ; that is

$$\ln Y_t = E(\ln Y_t) + \sigma_{\ln Y_t} Z_t$$

Therefore

$$EPP_t = 1 - \Pr\left(E[\ln Y_t] + \sigma_{\ln Y_t} Z_t \leq \ln Y_{up}\right),$$

or rearranging,

$$EPP_t = 1 - \Pr\left(Z_t \leq \frac{\ln Y_{up} - E[\ln Y_t]}{\sigma_{\ln Y_t}}\right) \quad (4)$$

To implement (4), we need to calculate $E[\ln Y_t]$ and $\sigma_{\ln Y_t}$. To do so, note that iterating (3) backwards implies that

$$E[\ln Y_t] = \ln Y_0 + \left(\mu - \frac{\sigma^2}{2}\right)t$$

and

$$\sigma_{\ln Y_t} = \sigma\sqrt{t}$$

Substituting these into (4) allows us to calculate the excess payment probability. Example 1 illustrates.

Example 1: Uncapped utility connection subsidy—estimating the excess payment probability (annual expenditure) using an analytical approach

The government is considering offering a utility provider a \$100 subsidy for every new connection made in certain rural areas in the next five years, with payment of the subsidy being made at the end of the year in which the connection occurs. The current annual rate at which new connections are being made is 10,000, but following the introduction of this scheme, this is expected to grow at a rate of 10 percent a year with a standard deviation of 20 percent a year; that is, $\mu = 0.1$ and $\sigma = 0.2$. Since $Y_t = 100$ multiplied by the number of new connections in year t , it follows that

$$E[\ln Y_t] = \ln(100 \times 10,000) + \left(0.1 - \frac{.04}{2}\right)t$$

$$\sigma_{\ln Y_t} = 0.2\sqrt{t}$$

Suppose we want to know for each of the five years of the scheme the probability of payments exceeding \$2 million. Then we need to apply equation (4) for $t = 1$ to 5. For $t = 1$, we first find that

$$\Pr\left(Z_1 \leq \frac{\ln Y_{up} - E[\ln Y_1]}{\sigma_{\ln Y_1}}\right) = \Pr(Z_1 \leq 3.07).$$

We can use Excel's "NORMSDIST" function (or the Z-distribution tables found in statistics texts) to calculate this probability. Substituting it into (4) and doing the same for the following four years, we get the following results.

Year	Excess payment probability
1	0.001
2	0.030
3	0.095
4	0.175
5	0.256

In the first year of the scheme, there is only a 0.1 percent (that is, 1 in 1000) chance that the subsidy payment will exceed \$2 million; by year 5 the probability has risen to approximately 26 percent.

3.4.2 Other payment distributions—The use of Monte Carlo simulation

In some circumstances, it might not be appropriate to assume that the relevant payment has either a normal or a lognormal distribution. For example, it may be the case that the government is

interested in total payments over the life of the scheme, rather than just payments in a single year. Then the variable of interest is

$$Y_t^{Total} = \sum_{h=1}^t Y_h.$$

Because the log of a sum of variables does not equal the sum of the logged variables, $\ln Y_t^{Total}$ does not have a normal distribution, so we cannot use (4). In this situation, the best way of proceeding is to generate the underlying probability distribution of the variable of interest and calculate the excess payment probability directly. For this purpose, we can use a method known as *Monte Carlo simulation*. In simple terms, this works by using a random number generator (such as RAND(.) in Excel) to create many alternative realizations of a variable, each of which is consistent with our information about the distribution (such as μ and σ). In this way, we build up a picture of the variable's entire probability distribution. For instance, suppose that Y_t is given by (3). Then a Monte Carlo simulation of Y_t simply involves using a random-number generator to produce a large number of values of Z_t , each of which is used in (3) to calculate a corresponding value of Y_t .⁷ Each of these represents one possible realization of Y_t ; together they provide a picture of the entire distribution of Y_t . The proportion of Y_t values exceeding Y_{up} is our estimate of the excess payment probability.

As an example of how this procedure works for more-complex situations, suppose that Y_t is given by (3) and that the variable of

⁷ The more trials, the more accurate the estimate. Working out how many trials to use requires balancing the time it takes a computer to generate a trial against the desire for smaller simulation errors. For suggestions on determining the appropriate number of trials, see Campbell et al (1997, p384), or Hull (2003, p413).

interest is Y_t^{Total} ; we wish to know the probability of this exceeding some upper bound Y_{up}^{Total} . Using Monte Carlo simulation, we proceed as follows:

- i. Use a random number generator to produce a large number of sample paths for Z (for example, 10,000).
- ii. For each realization of Z_t , use (3) to calculate the corresponding value of Y_t . This produces a large number (for example, 10,000) of sample paths for Y .
- iii. For each Y path, calculate the value of Y_t^{Total} corresponding to every Y_t . This produces a large number (for example, 10,000) of sample paths for Y^{Total} .
- iv. For each t , calculate the proportion of Y^{Total} values that exceed Y_{up}^{Total} . This is the excess payment probability for that t .

Example 2 illustrates.

Example 2: Uncapped utility connection subsidy—estimating the excess payment probability (total expenditure) using Monte Carlo simulation

Continue with the situation described in Example 1, but now suppose we want to know the probability that total connection subsidy payments over the entire life of the scheme exceed \$10 million. Because the sum of payments does not have a normal distribution, we use Monte Carlo simulation, following steps (i)–(iv) above.

The second and third columns in the table below are two randomly generated sample paths for Z (step (i)). Column four substitutes each realization of Z into equation (3):

$$Y_t = Y_{t-1} \exp \left[\left(\mu - \frac{\sigma^2}{2} \right) + \sigma Z_t \right]$$

to obtain the values of Y_t corresponding to each realization of Z in the first sample path (step (ii)). For example, given $Z(1) = 0.40$ at date 1, and $Y_0 = 100(10,000)$, the corresponding value of Y_1 is, in millions of dollars,

$$Y_1(1) = (1) \exp \left[\left(0.1 - \frac{.04}{2} \right) + 0.2(0.40) \right] = 1.174$$

Column five then calculates the value of Y_t^{Total} corresponding to each of these realizations of Y_t (step (iii)). The next two columns repeat these two steps for the second sample path of Z . Finally, the last column calculates the proportion of outcomes for which Y_t^{Total} exceeds \$10 million (step (iv)). For example, in year 1, neither possible value for Y_t^{Total} exceeds \$10 million, so the estimated excess payment probability for that year is zero. In year 5, by contrast, sample path 2 still yields a Y_t^{Total} figure less than \$10 million, but the sample-path-1 figure is \$10.060 million. As the two paths are equally likely to occur, the estimated excess payment probability in year 5 is 0.5.

Year	Z(1)	Z(2)	Y(1)	Y^{Total} (1)	Y(2)	Y^{Total} (2)	EPP
1	0.40	-0.17	1.174	1.174	1.048	1.048	0
2	0.68	-0.17	1.456	2.629	1.096	2.144	0
3	0.98	-1.52	1.918	4.548	0.876	3.021	0
4	1.33	-0.97	2.711	7.258	0.781	3.802	0
5	-0.23	0.08	2.802	10.060	0.861	4.663	.5

Note: in millions of dollars, except for Z-values and EPP.

Of course, an estimate based on just two trials is likely to be wildly inaccurate, but the large number of trials needed to produce a reliable estimate does not change the procedure. Extending the simulation to 10,000 trials, we get an estimate of the probability that expenditure exceeds \$10 million in year 5 of approximately 7 percent.

3.4.3 Other Itô processes—Monte Carlo simulation again

In the above example, Y follows a geometric brownian motion and thus is given by (3), but the variable whose risk we wish to determine is some transformation of Y , and thus is not a geometric brownian motion. As we have seen, this problem is resolved by using Monte Carlo simulation. In other situations, Y itself may not satisfy (3). This too poses no problem for Monte Carlo simulation. If, for example, Y follows the more general Itô Process given by (1), then so long as we have some mechanism for specifying values of the drift and diffusion parameters at each date, we simply use those date-specific values (rather than the constant μ and σ values assumed above) and proceed exactly as before. We return to this issue in section 6.1.

3.5 Cash-flow-at-risk

The excess-payment-probability measure estimates the probability of payments exceeding some fixed value. An alternative risk measure estimates the maximum payment consistent with some fixed probability.⁸ This measure, known as cash-flow-at-risk (CAR), seeks to estimate the payment associated with the α percentile of the distribution. For example, if $\alpha = 95$, then the cash-flow-at-risk is the biggest payment expected with a 95 percent degree of confidence; that is, the actual payment is expected to be larger than the cash-flow-at-risk only $(1 - \alpha)$ percent = 5 percent of the time. Thus, cash-flow-at-risk provides an approximate “upper bound” estimate of the payment. All else equal, a scheme with a high cash-flow-at-risk is riskier than a scheme with a low cash-flow-at-risk.

⁸ This measure is originally due to Baumol (1963). More recently, it has been rediscovered and given a new name—Value (or Cashflow)-at-Risk. See, for example, Dowd (1998).

As with excess payment probability, cash-flow-at-risk (CAR) can be found using a simple formula if Y has a normal distribution. In this case, the α percentile of the distribution is given by a known number of standard deviations over and above the mean. That is

$$CAR = E[\text{payment}] + \sigma(\text{payment})\alpha^* \quad (5)$$

where the value of α^* corresponding to each confidence level α can be found in the standard normal distribution tables or by using the Excel function `NORMSINV(α)`; for example, if $\alpha = 95$, $\alpha^* = 1.645$.

Unfortunately, equation (5) is not likely to be very helpful for assessing the risk of most output-based-payment schemes, except as a rough approximation. Not only is Y unlikely to be normal, but using $\ln Y$ in its place is no help in this case.⁹ As a result, the cash-flow-at-risk associated with output-based aid schemes must generally be calculated using Monte Carlo simulation. This consists of, first, generating the simulated probability distribution for Y as before and, second, identifying the cash-flow-at-risk (see Example 3).

Example 3: Uncapped utility-connection subsidy—estimating cash-flow-at-risk using Monte Carlo simulation

Consider again the utility-connection scheme described in Examples 1 and 2. Suppose we wish to know, with 66-percent confidence, the maximum payment the government is likely to have to make in each year of the scheme (that is, the payment

⁹ To see why, suppose we calculated

$$\alpha\text{-maximum (log) payment} = E[\ln \text{payment}] + \sigma(\ln \text{payment})\alpha^*$$

and then tried to back out the α -maximum payment. Unfortunately, there is no simple way of accomplishing the last step because, for example, $E[\ln \text{payment}] \neq \ln E[\text{payment}]$, so $\exp\{E[\ln \text{payment}]\} \neq E[\text{payment}]$. More generally, α -maximum (log) payment is not the same as the log of α -maximum payment; we need the latter, but equation (5) can give us only the former.

amount the government can be 66-percent sure it will not have to exceed).

In the table below, the second through fourth columns contain three sample paths for the standard normal random variable Z . Columns five through seven use (3) to calculate the corresponding sample paths for Y . With only three such paths, the 66th percentile is equal to the second-highest realization of Y . This appears in the final column.

Year	Z(1)	Z(2)	Z(3)	Y(1)	Y(2)	Y(3)	CAR (\$ million)
1	0.96	0.93	-1.90	1.312	1.306	0.741	1.306
2	0.73	-1.75	-0.16	1.643	0.996	0.777	0.996
3	1.49	-0.16	1.56	2.398	1.045	1.150	1.150
4	-0.86	0.99	0.14	2.185	1.380	1.282	1.380
5	0.34	-0.35	1.65	2.535	1.393	1.933	1.933

For example, at date 3, the subsidy payment exceeds \$1.15 million with 0.34 probability, so CAR = \$1.15 million.

In practice, obviously, governments are likely to be concerned with percentiles higher than the 66th, such as the 90th, the 95th, or the 99th. We've used the 66th percentile confidence level here only because it allows us to illustrate Monte Carlo simulation for calculating cash-flow-at-risk using just three trials. The same technique can be used to estimate results for any value of α . For example, suppose we are interested in estimating cash-flow-at-risk for $\alpha = 0.99$ in the fifth year. In a Monte Carlo simulation with 10,000 trials, the 9,900th largest payment provides an estimate of cash-flow-at-risk at the 99th percentile. Using such a Monte Carlo simulation, we obtain an estimate of \$4.2 million. That is, payments are lower than \$4.2 million in year 5 in the simulation 99 percent of the time.


3.6 Estimating the entire probability distribution

The standard deviation of payments, the excess-payment probability, and cash-flow-at-risk are all single-number measures of the fiscal risks created by output-based-payment schemes. Each provides valuable information, and for some purposes one of them may provide as much information as is needed. If a government is concerned only to limit the probability of overspending on the scheme, for example, the excess-payment probability gives it all the information it needs. No single number can fully describe the risks the government faces, however, and it may often be useful to present more than a single number. Indeed, it may be helpful to present a picture of the entire probability distribution, such as a histogram showing an estimate of the probability of payments falling in each of several bands (or “bins”).

When the probability distribution is known, it is usually possible to describe the risks fully, and depict them in a histogram, on the basis of just two or three parameters. For example, if the payments are normally distributed, it is sufficient to specify the mean and the standard deviation. In other cases, it isn't possible (or at least isn't easy) to specify the probability distribution in this way. As before, however, we can still estimate the probability distribution using Monte Carlo simulation.

In Example 3, we estimated that the cash-flow-at-risk at the 99th percentile in the fifth year was \$4.2 million. More revealing is a picture of the entire probability distribution, as shown in Figure 1. It shows the government will most probably have to make a payment of between \$1 and \$1.5 million (the bar labeled “1.5” shows payments in this range). The frequency of such payments in the simulation is 3,151 out of a possible 10,000, so the probability is approximately 32 percent. It also shows that there is only an estimated 36 out of 10,000 probability of the government having to make payments greater than \$5 million; a 67 (= 36 + 31) chance out of 10,000 of payments of more than \$4.5 million; and so on.

Figure 1: Estimated frequency distribution of payments in year five



Note: the bin on the far left, labeled “0” shows the frequency in the simulation of payments of zero or less (in this case, zero). The next, labeled “0.5,” shows the number of payments between 0 and \$0.5 million (in this case 75). The rightmost bin, labeled “More,” shows the frequency of payments greater than \$5 million (in this case 36).

Table 2 summarizes the discussion of four measures of risk discussed here.

Table 2 Risk Measures for Output-Based Subsidy Schemes

Measure	Description	Advantages	Disadvantages
Standard deviation	Standard deviation of annual change in payments	Provides government with a single number summarizing how variable payments are	Doesn't distinguish between upside and downside risk
Excess-payment probability	Probability that subsidy payments exceed X	Provides government with a single number that helps determine whether risk to government's fiscal position is significant	Doesn't offer much information on the probabilities of other payments
Cash-flow-at-risk	Maximum payment with $\alpha\%$ probability	Provide government with a single number that helps determine whether risk to government's fiscal position is significant	Doesn't offer much information on other possible payments; may be mistaken for maximum possible payment
Frequency distribution of payments	Probability of payments in each of several intervals	Provides government with a picture of the entire range of possible outcomes	The information requires a graph or table to convey; it is not succinct

All these approaches—measuring excess payment probability and cash-flow-at-risk and presenting a picture of the whole probability distribution—are reasonable ways of describing the risks of payment schemes. However, as we shall see in the next section, the risks measured in these ways are not necessarily the risks that necessitate a premium in estimating the cost of a scheme. It is to this latter issue that we now turn.

4 General principles for valuing output-based schemes

In this section, we outline the general principles underlying the methods that can be used to estimate the *cost* of various kinds of payment scheme. In the next section, we apply these methods to the schemes previously described in section 2.

As before, let Y_t denote the date t cash payment made by the government to a private investor in return for the provision of some agreed service. What we would like to be able to do is calculate the current dollar value of this obligation.

One option is simply to discount the expected value of Y_t at the riskless rate of interest. This approach will be good enough for certain purposes. But it ignores the price of bearing risk and may generate a poor estimate of the value of some obligations. For large, long-term, risky commitments, the government may want to employ an approach to valuation that incorporates the price of bearing risk.

The valuation approach described below estimates the cost of the scheme as if the cash flows generated by the scheme were available for trading in financial markets. That is, we estimate the market value of the scheme in the context of all available assets and securities. Alternatively, we could estimate this cost in the context of the provider's portfolio alone. The principles underlying these two approaches are the same, but they differ in one fundamental way. In the market approach, the scheme's payoff structure is implicitly already available to financial market investors, so its cost is a function of current market prices; we discuss this point further below equation (7). In the portfolio-specific approach, the scheme's payoff structure is explicitly a new asset whose cost depends on the additional "opportunities" it creates. One disadvantage of the latter

approach is that introduction of a new scheme requires revaluation of the entire portfolio.

4.1 Adjusting the discount rate

To value an output-based scheme, we need some way of converting each of the unknown future payments into a current (and therefore known) dollar equivalent. That is, how many birds in the hand correspond to two in the bush? One popular method for addressing this issue is known as the *risk-adjusted-discount-rate* approach and consists of the following three steps:

Estimate the expected value (that is, the mean) $E[Y_t]$ of each future payment.

Adjust each expected payment using a discount rate k commensurate with the risk of the payment stream.

Add all these adjusted payments together.

The logic underlying this procedure is straightforward. The first step estimates the expected (or “average”) value of each future payment. The second step converts each of these values into a current-date equivalent by working out what sum of money today would, if invested at the risk-adjusted-discount-rate, grow to equal the expected (“average”) future value. Finally, the third step aggregates all these individual current-dollar values into a single, total, value.

The risk-adjusted discount rate method is an example of the well-known *discounted-cash-flow* approach to valuation. The discount rate k “translates” each expected future dollar payment $E[Y_t]$ into an equivalent number of current dollars by calculating what must be invested today at expected return k in order to generate the amount $E[Y_t]$ at date t . For example, if the government expects to make a payment of \$100,000 in one year and k is currently equal to

0.1, then $\$(100,000/1.1) = \$90,909$ is the amount that must be invested today at 10 percent in order to generate \$100,000 in one year. In this sense, \$90,909 is the *present value* of \$100,000 in one year.

To progress further, we need to make all this a bit more precise. If the payment scheme continues until date T , the risk-adjusted discount rate approach states that its total cost C in current dollars is

$$C = \frac{\sum_{t=1}^T E[Y_t]}{(1+k)^t} \quad (6a)$$

where k is the required rate of return on the payment stream Y (that is, the risk-adjusted discount rate).

Equation (6a) assumes that payments occur at discrete points in time. For some purposes, it is mathematically more convenient to assume that these occur continuously, in which case (6a) becomes

$$C = \int_0^T E[Y_t] e^{-kt} dt \quad (6b)$$

To implement (6a) or (6b), we obviously need estimates of both the expected payment stream $\{E[Y_t]\}$ and the discount rate k . Again, as noted above in section 3, the precise method for estimating $\{E[Y_t]\}$ may vary from case to case, and may require expert advice. Our primary focus is on the core issues surrounding estimation of the discount rate k .¹⁰

In general, the discount rate has two components

$$k = \text{riskless rate of interest} + \text{risk premium}$$

¹⁰ For further reading, see, for example, Brealey and Myers (2000) or, at a more advanced level, Cochrane (2001).

$$\equiv r + \lambda$$

In practice, the riskless rate r is usually proxied by a short-term government bond. To estimate the risk premium λ , we need a model that provides a mechanism for estimating the price of bearing risk. Of these, the most popular and widely used is the celebrated Capital Asset Pricing Model (CAPM).¹¹ This states that

$$\lambda = \beta_m \{E[R_m - r]\} \quad (7)$$

where

$$\beta_m = \frac{\text{cov}(\%Y, R_m)}{\sigma_m^2}$$

is the “beta” of the payment stream Y_t

measuring the extent to which percentage changes in Y ($\%Y$) are linearly related to returns on the market portfolio (R_m).

$\text{cov}(\%Y, R_m)$ = the covariance between percentage changes in Y and R_m

σ_m^2 = the variance of R_m .

Equation (7) states that the risk premium associated with the scheme’s payment stream is a multiple of the market risk premium $\{E[R_m] - r\}$, where the multiple equals the proportion of the market’s risk that is contained in the payment stream. This result is an outgrowth of modern portfolio theory, which stresses the importance of diversification for managing the risk of asset portfolios. By combining assets with different risk characteristics, the potential for negative returns on one asset can be partially offset by positive returns on other assets, so the risk of a portfolio is

¹¹ The original papers on the CAPM are Sharpe (1964) and Lintner (1965). For a textbook exposition, see Brealey and Myers (2000). In section 6, we discuss other possible models for estimating the risk premium.

generally less than the risk of a single asset in isolation. The risk that can be eliminated in this fashion is known as diversifiable risk; the risk that cannot be eliminated by diversification is known as systematic risk and reflects the extent to which returns on a single asset are correlated with returns on a portfolio of assets. As investors can, subject to limitations imposed by transactions costs, freely choose their portfolios to eliminate diversifiable risk, they cannot expect to receive a premium for continuing to bear such risk, so an asset's risk premium should reflect only its systematic risk. And because *all* diversifiable risk is removed from the portfolio that contains all assets (the so-called market portfolio), an asset's systematic risk is measured by the correlation between its returns and returns on the market portfolio; in equation (7), the extent of this correlation is measured by "beta."

Intuitively, a high beta means that payments tend to be low when the returns on other assets are also low, so these payments do not offset the recipient's other income streams and thus have relatively little diversification value. As a result, they have high risk and are thus discounted heavily. Similarly, a low beta means that payments are more likely to be high when the returns on other assets are low, so these payments hedge the recipient's other income streams and thus have high diversification value. Consequently, they have low risk and the discount rate is correspondingly low.

Some readers may find it puzzling that a high-beta payments stream should be discounted heavily, and thus implicitly entered on the government's books at a low cost. After all, if governments (and their citizens) are averse to risk, then it might seem reasonable that a high-beta scheme should have a *higher* cost than a low-beta scheme. This can be resolved by recalling that, from the government's perspective, a scheme commits it to *making*, rather than receiving, payments; that is, it is akin to a negative asset, or a liability, in terms of the portfolio theory outlined above. That is, a high-beta payments stream is one that commits the government to making high payments when it can most afford to do so (when

other assets in the government's portfolio are doing well) and to making low payments when conditions are less favorable (when other assets are doing poorly). In short, a high-beta scheme smoothes the government's overall financial commitments and thus is more attractive than a low-beta scheme.¹²

4.2 Estimating certainty equivalents

The risk-adjusted discount rate approach to valuation incorporates risk by adjusting the discount rate. An alternative approach, known as the *certainty-equivalent* method, incorporates the risk adjustment directly in the expected payments. Under this method, we estimate the expected payments using not the true probability distribution of payments, but rather an artificial probability distribution that replaces the time- t expected growth rate μ_t with the time- t risk-adjusted expected growth rate $(\mu_t - \lambda)$. That is, wherever we would have used the true expected growth rate μ_t in calculating $E[Y_t]$, instead use the risk-adjusted growth rate $(\mu_t - \lambda)$ to calculate the certainty-equivalent payment $E^*[Y_t]$. For example, the expected value of next year's payment is

$$E[Y_1] = Y_0(1 + \mu_1),$$

while the certainty-equivalent analogue is

$$E^*[Y_1] = Y_0(1 + \mu_1 - \lambda)$$

The expected-value transformation $E[Y]$ contains no adjustment for risk, so the appropriate discount rate is a risk-adjusted one. By

¹² This discussion is couched in terms of the CAPM risk premium, but it applies generally. For any income stream that has a positive correlation with risk factors, the corresponding payments stream must have a negative correlation with these factors, since each possible payment by the government is simply the negative of the income to the recipient.

contrast, the certainty-equivalent transformation $E^*[Y]$ is already risk adjusted, so the appropriate discount rate is the riskless rate r .¹³

The three steps underlying the certainty-equivalent approach are, therefore:

Estimate the certainty-equivalent value $E^*[Y_t]$ of each future payment.

Discount each certainty-equivalent payment at the riskless rate of interest r .

Add all these adjusted payments together.

Mathematically, we can express this as

$$C = \frac{\sum_{t=1}^T E^*[Y_t]}{(1+r)^t} \quad (8a)$$

or, with continuous payments and compounding

$$C = \int_0^T E^*[Y_t] e^{-rt} dt \quad (8b)$$

4.3 Equivalence in principle of the two methods

The risk-adjusted discount rate and certainty-equivalent approaches are in fact equivalent. To see this in a simple way, let's compare (6b) and (8b) assuming that the payments follow the continuous geometric brownian motion process described in equation (2). That is, let

¹³ For an introduction to certainty-equivalent valuation, see Brealey and Myers (2000). More advanced treatments appear in Cox and Ross (1976) and Cox, Ingersoll, and Ross (1985).

$$\frac{dY}{Y} = \mu dt + \sigma dz$$

where μ and σ are constants. Since expected payment growth is a constant μ , the expected payment is

$$E[Y_t] = Y_0 e^{\mu t} \quad (9)$$

where Y_0 is the current value of Y . Then the risk-adjusted discount rate method implies

$$C = Y_0 \int_0^T e^{\mu t} e^{-(r+\lambda)t} dt \quad (10)$$

As described above, replacing μ with $(\mu - \lambda)$ in (9) gives us the certainty-equivalent payment, so

$$E^*[Y_t] = Y_0 e^{(\mu - \lambda)t}$$

and (8b) becomes

$$C = Y_0 \int_0^T e^{(\mu - \lambda)t} e^{-rt} dt$$

Using the properties of the exponential function, this can be rewritten as

$$C = Y_0 \int_0^T e^{\mu t} e^{-(r+\lambda)t} dt$$

which is the same as (10).¹⁴

¹⁴ For the more general continuous-time Itô process in equation (1), the mathematics become more complex, but equivalence continues to hold; see Cox, Ingersoll, and Ross (1985). In discrete time, the equivalence is only approximate, since $(1 + \mu - \lambda)/(1 + r)$ is not generally equal to $(1 + \mu)/(1 + r + \lambda)$. However, the degree of difference is small so long as r , μ , and λ are small.

4.4 Superiority in practice of the certainty-equivalent method

The equivalence of the risk-adjusted discount rate and certainty-equivalent approaches suggests we should be indifferent between these two methods. However, it turns out that it is often easier to use the certainty-equivalent approach for valuing the government's liability under output-based aid schemes. The reason for this is that the payment streams associated with these schemes often depend on the realization of other, more fundamental, underlying variables. For example, with a connection subsidy

$$Y_t = sX_t$$

where s is the dollar subsidy paid for each new connection and X is the number of new connections. If the annual connection subsidy is capped,

$$Y_t = s \max\{X_t, \bar{X}_t\}$$

where \bar{X} is the maximum number of new connections that the government will subsidize in a year.

In such a case, among others, the payment stream Y generated by a subsidy scheme is some function of an underlying variable (in this case, the number of new connections). This creates problems for the risk-adjusted-discount-rate approach, because the implementation of (6) requires estimation of a discount rate that is adjusted for the risk of the payment stream Y . For simple payment streams whose risk characteristics remain constant over time, we can do this by using historical data to estimate the parameters of the CAPM (or some other asset pricing model) and thus calculate the risk premium λ . However, payment streams of the sort described above will generally not have constant risk, even if the underlying

variable does have this property.¹⁵ Thus, we have no feasible way of estimating the appropriate discount rate, a problem that is compounded if there are multiple underlying variables.

However, it turns out that this problem does not arise with the certainty-equivalent approach, even when the payment depends on multiple underlying variables: the calculation of the certainty-equivalent payments requires only the risk-adjusted growth rates of the *underlying* variables, and *not* the risk-adjusted growth rate of Y .¹⁶ As a result, the certainty-equivalent approach requires only the risk premiums λ_i associated with each of the underlying variables X_i , each of which we can (at least in principle) estimate using standard methods.

To summarize, suppose the payment stream Y_t associated with a payment scheme depends on one or more underlying random variables X_i , each of which follows a geometric brownian motion with drift μ_i and volatility σ_i . Then the value of this payment stream can be found by undertaking the following procedure:

Step 1. Estimate the risk premium λ_i associated with each underlying variable.

Step 2. Use these risk premiums to calculate the certainty-equivalent value $E^*[Y_t]$ of each future payment Y_t . That is, calculate the expected value of each future payment Y_t using the *risk-adjusted* expected growth rates μ_i^* of each underlying variable X_i , where

¹⁵ For example, suppose $Y = f(X)$ for some arbitrary function $f(\cdot)$ and underlying variable X . Then the CAPM risk premium for Y is proportional to $\text{cov}(\%Y, R_m) = \text{cov}(\%f(X), R_m)$, while the risk premium for X is proportional to $\text{cov}(\%X, R_m)$. Even if the latter is constant over time, the former will generally not be.

¹⁶ See Cox, Ingersoll, and Ross (1985).

$$\mu_i^* = \mu_i - \lambda_i.$$

Step 3. Multiply each $E^*[Y_t]$ by e^{-rt} to obtain the present value of each payment.

Step 4. Add all the payment present values $E^*[Y_t]e^{-rt}$ together to get the total cost of the scheme.

In the next section, we show how these four steps can be applied to a variety of schemes.

5 Valuation applications

5.1 Uncapped annual-output schemes

Of the types of schemes discussed in section 2, uncapped annual output schemes are the simplest to analyze. Example 4 illustrates the valuation of the government's liability under one form of such a scheme.

Example 4: Uncapped utility connection subsidy—valuing the subsidy using an analytical approach

Consider again the situation described in Example 1 and let X_t denote the number of new connections in year t . Suppose also that growth in new connections is estimated to have a 0.5 correlation with market returns, the riskless interest rate is 5 percent, and the excess market return (over and above the riskless rate) has a mean of 7 percent and a standard deviation of 20 percent.

To calculate the cost of this subsidy, we first need the risk premium associated with the risk in the rate of growth of connections (in other words, the price of connections-growth risk)— λ_x . For this purpose, we use the CAPM—see equation (7). To calculate β_m , we use the fact that $\text{cov}(x, y) = \rho(x, y)\sigma_x\sigma_y$, where $\rho(x, y)$ is the linear correlation between x and y . Then

$$\beta_m = \frac{\rho(\%X, R_m)\sigma\sigma_m}{\sigma_m^2} = \frac{\rho(\%X, R_m)\sigma}{\sigma_m} = \frac{0.5(0.2)}{0.2} = 0.5$$

so that

$$\lambda_x = \beta_m \{E[R_m - r]\} = 0.5\{0.07\} = 0.035$$

and, therefore

$$\mu^* = \mu - \lambda_x = 0.1 - 0.035 = 0.065$$

Because this scheme is relatively simple, we can value the government's liability using a standard formula for valuing a growing annuity. Specifically, the present value V of a T -year annuity that pays an amount A in year 0 and grows thereafter at a rate of g per year is given by

$$V = A \left(\frac{1+g}{g-r} \right) \left[\left(\frac{1+g}{1+r} \right)^T - 1 \right]$$

where r is the discount rate.

Here, we assume that the growth g rate is μ^* and then discount at the riskless rate. Noting that $A = sX_0 = 1$ million and that $T = 5$, we find that the value of the uncapped annual subsidy is \$5.218 million.

The complexities of real schemes often mean that applying a simple formula for a growing annuity isn't possible. Instead, it is necessary to calculate the value of each payment separately and then add the payments together. To illustrate this procedure, we apply it to the example above.¹⁷ In the first column of the table below, we calculate the certainty-equivalent payment for each year of the scheme. Column two gives the discount factor corresponding to each payment. Column three multiplies the first two columns together to get the present value of each future payment. Finally, each of these present values is added together to get the total cost—given in the bottom right corner.

¹⁷ The analytical simplicity of the scheme also means that using a risk-adjusted discount rate would be as easy as estimating certainty equivalents. We choose to illustrate the latter method here, because it is simpler in other cases.

Year	E*[Y] \$m	Discount factor	Value \$m
	[1]	[2]	[3]
1	1.065	0.95	1.014
2	1.134	0.91	1.029
3	1.208	0.86	1.043
4	1.286	0.82	1.058
5	1.370	0.78	1.073
Total cost (C_u^a)			5.218

With an expected growth rate of 10 percent and a risk premium of 3.5 percent, the certainty-equivalent growth rate is 6.5 percent. Thus, the certainty equivalent payment for the first year is, in million of dollars, $1 \times 1.065 = 1.065$. Discounting this back one year at 5 percent gives a present value of $1.065/1.05 = 1.014$. Continuing in this manner for years 2–5 produces the rest of the table.

5.2 Capped annual-output schemes

When the annual payment is capped, the year- t payment is

$$Y_t = \min\{sX_t, s\bar{X}\},$$

where \bar{X} is the maximum output level on which the government will pay s ; that is, the maximum payment is $s\bar{X}$. This scheme is more complicated than the uncapped version, but payments can still be valued. To see how, note that by adding and subtracting sX_t , Y_t can be rewritten as

$$Y_t = sX_t + \min\{0, s\bar{X} - sX_t\}.$$

Noting that $\min\{a, b\} = -\max\{-a, -b\}$, this can be rewritten as

$$Y_t = sX_t - \max\{0, sX_t - s\bar{X}\}. \quad (11)$$

The first term on the right-hand side of (11) is the payment on an uncapped scheme; the second term is the payoff function for a t -year call option on sX with an exercise price of $s\bar{X}$. Intuitively, the cap gives the government the right to “call” (that is, retain) any payments above $s\bar{X}$, so the uncapped payment is reduced by an amount equal to the benefit the government gets from the option. Thus, the value of a capped annual-output scheme C_c^a is the value of the equivalent uncapped scheme C_u^a less the value of T call options on payments exceeding the cap:

$$C_c^a = C_u^a - \sum_{t=1}^T V_t,$$

where V_t is the value of an option giving the holder the right to sX_t in exchange for a payment of $s\bar{X}$. The value of each option can be found using a modified version of the Black-Scholes equation¹⁸

$$V_t = sX_0 e^{(\mu^* - r)t} N(d_1) - s\bar{X} e^{-rt} N(d_2) \quad (12)$$

where $N(\cdot)$ is the distribution function for the standard normal random variable (NORMSDIST(\cdot) in Excel), and

$$d_1 = \frac{\ln\left(s \frac{X_0}{\bar{X}}\right) + \left(\mu^* + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}}$$

$$d_2 = d_1 - \sigma\sqrt{t}.$$

Example 5 illustrates.

¹⁸ The modification comes about because the Black-Scholes equation applies to the specific situation where $\mu = r + \lambda$, which is not the case here. For further reading on the Black-Scholes equation, see, for example, Hull (2003).

Example 5: Annually capped utility connections subsidy—valuing the subsidy using an analytical approach

In the previous example, suppose that annual payments are capped at $s\bar{X} = \$1.2$ million. Using equation (12), we get (in millions of dollars)

$$V_1 = .037$$

$$V_2 = .093$$

$$V_3 = .150$$

$$V_4 = .205$$

$$V_5 = .259$$

(Note that the option payoff increases monotonically with payment date. This is a standard property of options: long-dated options offer more opportunity for profit than otherwise-equivalent short-dated options.)

Therefore

$$\sum_{t=1}^T V_t = .743$$

and

$$C_c^a = C_u^a - \sum_{t=1}^T V_t = 5.218 - 0.743 = 4.475$$

That is, the annual payment cap reduces the cost of the subsidy by \$0.743 million.

5.3 Schemes with multiple payment rates

Where a scheme has capped annual expenditure, there are effectively two subsidy rates— s and zero. More generally, however, a scheme could offer multiple payment rates. For example, it might pay the rate s_1 up to some threshold volume of output \bar{X}_1 , then the

rate s_2 on output \bar{X}_1 and some higher level \bar{X}_2 , and then zero thereafter (as in the case of the shadow toll road discussed in section 2). In this case, the payment is

$$Y_t = \min\{s_1 X_t, s_1 \bar{X}_1 + s_2 (X_t - \bar{X}_1), s_1 \bar{X}_1 + s_2 (\bar{X}_2 - \bar{X}_1)\}$$

This payment structure is more complex than that offered by the simple cap, and thus offers additional challenges. In particular, simple valuation formulas are usually unavailable for schemes of this type. Fortunately, we can still obtain a cost estimate by combining the Monte Carlo simulation approach outlined in section 3.4.2 with the certainty-equivalent valuation method. In effect, we repeat the procedure detailed in section 3.4.2, but replace μ with $\mu - \lambda_x$ in generating the sample paths for the variable being simulated.

Specifically, we undertake the following steps

- (i) Use the information we have about the certainty-equivalent distribution of X to simulate a large number of possible time series paths for that variable; that is, simulate paths for X assuming that this variable has an expected growth rate of $\mu - \lambda_x$.
- (ii) Calculate the payment payoff Y corresponding to each realization of X ; this generates a large number of possible time series paths for Y .
- (iii) At each date t , calculate the average of all the sample realizations of Y ; this generates an estimate of $E^*[Y_t]$.

- (iv) Discount each $E^*[Y_t]$ by the date- t riskless discount factor and then add all these present values together to obtain a total cost estimate.¹⁹

Example 6: Scheme with multiple payment rates—valuation using Monte Carlo simulation

Take the information from Example 4, but suppose that output is traffic on a shadow toll road. Suppose that the initial annual volume is 1 million vehicles and that the annual schedule of shadow tolls per million vehicles is as follows:

For $X_t \leq 1.2$, $s_1 = \$1$

For $X_t \in (1.2, 1.5)$, $s_2 = \$0.5$

For $X_t > 1.5$, $s_3 = 0$.

With this schedule, total expenditure is effectively capped at \$1.35 million ($=1.2(1) + (1.5-1.2)(0.5)$).

To illustrate the use of Monte Carlo simulation in estimating the value of such a scheme, we create two sample paths for X_t and then follow the procedure described above. In the table below, we first use a random number generator to produce two simulated paths for the standard normal random variable Z ; these appear in the second and third columns of the table. The next two columns substitute each realization of Z into the certainty-equivalent form of equation (3)

$$X_t = X_{t-1} \exp \left[\left(\mu - \lambda - \frac{\sigma^2}{2} \right) + \sigma Z \right] \quad (3a)$$

¹⁹ The use of Monte Carlo simulation for valuation purposes was first suggested by Boyle (1977). For further reading, see Hull (2003) or Dowd (1998).

to obtain corresponding simulated paths for the number of connections X_t (assuming that the expected growth rate is $\mu - \lambda$). For example, given $Z(1) = 0.399$ at date 1, the corresponding value of X , in millions of dollars, is

$$X_1(1) = \exp \left[\left(0.1 - 0.035 - \frac{(0.2)^2}{2} \right) (1) + 0.2\sqrt{1}(0.399) \right] = 1.133$$

The columns titled $Y(1)$ and $Y(2)$ use

$$Y_t = \min \left\{ s_1 X_t, s_1 \bar{X}_1 + s_2 (X_t - \bar{X}_1), s_1 \bar{X}_1 + s_2 (\bar{X}_2 - \bar{X}_1) \right\}$$

to obtain sample paths for Y_t . For example, given $X_1(1) = 1.133$ and the schedule of payment rates set out above, the corresponding value of Y is

$$Y_1(1) = \min \left\{ \begin{array}{l} (1)(1.133\text{m}), \\ (1)(1.2\text{m}) + (0.5)(1.133 - 1.2\text{m}), \\ (1)(1.2\text{m}) + (0.5)(1.5\text{m} - 1.2\text{m}) \end{array} \right\} = 1.133\text{m}$$

The average of $Y(1)$ and $Y(1)$ is then calculated to obtain $E^*[Y]$. After discounting at the riskless rate, the total estimated cost of the scheme (using just two trials) is \$4.984 million.

Year	Z(1)	Z(2)	X(1)	X(2)	Y(1)	Y(2)	E*(Y)	Discount factor	Value
1	0.40	0.41	1.133	1.134	1.133	1.134	1.134	0.95	1.078
2	-0.71	-0.26	1.028	1.127	1.028	1.127	1.077	0.90	0.975
3	-0.41	0.49	0.991	1.299	0.991	1.250	1.120	0.86	0.964
4	-0.12	1.76	1.013	1.933	1.013	1.350	1.182	0.82	0.967
5	0.75	-0.59	1.232	1.798	1.216	1.350	1.283	0.78	0.999
Total									4.984

Using the same procedure but running 10,000 trials, we get a more reliable estimate of \$4.7 million.

5.4 Total cumulative cap

The Monte Carlo approach is also necessary when total (rather than annual) expenditure on the scheme is capped at some level \bar{C} . In this case, payments are given by

$$Y_t = \min \left\{ sX_t, \max \left\{ 0, \bar{C} - \sum_{h=1}^{t-1} sX_h \right\} \right\}.$$

By adding and subtracting sX_t and recalling that $\min\{a, b\} = -\max\{-a, -b\}$, we get

$$Y_t = sX_t - \max \left\{ 0, sX_t - \max \left\{ 0, \bar{C} - \sum_{h=1}^{t-1} sX_h \right\} \right\},$$

which states that the year- t payment on the capped scheme is equal to the uncapped payment less the payoff on a call option giving the rights to sX_t in exchange for an exercise price of $\max \left\{ 0, \bar{C} - \sum_{h=1}^{t-1} sX_h \right\}$. There is no simple formula for calculating the value of an option with such a complex exercise price, and thus no easy way to obtain the certainty-equivalent payment for each year

the scheme operates. However, we can again use the Monte Carlo simulation approach, as Example 7 shows.

Example 7: Utility connections subsidy with cumulative cap—valuing the subsidy using Monte Carlo simulation

Return to the subsidy described in Example 1, and suppose that total payments are capped at $\bar{C} = \$6$ million. Then payments in each year equal $\$100(X_t)$ until $\sum_{h=1}^t 100(X_h) \geq 6,000,000$, at which point payments become zero. The second and third columns of the table below are simulated paths for Z . The next two columns substitute each realization of Z into equation (3a) to obtain corresponding simulated paths for the uncapped payoff sX_t , assuming the expected growth rate is $\mu - \lambda_x = 0.065$. The columns titled $Y(1)$ and $Y(2)$ use

$$Y_t = \min \left\{ sX_t, \max \left\{ 0, 6 - \sum_{h=1}^{t-1} sX_h \right\} \right\}$$

to obtain sample paths for Y_t . For example, in path (2), the total expenditure at the end of year 3 is \$5.768 million. This leaves only \$0.232 available for years 4 and 5, and this is less than the simulated \$2.705 million for year 4, so $Y_4(2) = \$0.232$ million, and $Y_5(2) = 0$. The average of $Y(1)$ and $Y(2)$ is then calculated to obtain $E^*[Y]$. After discounting, the total cost of the scheme is \$5.247 million.

Year	Z(1)	Z(2)	sX(1)	sX(2)	Y(1)	Y(2)	E*(Y)	Discount factor	Value
1	0.76	0.39	1.218	1.132	1.218	1.132	1.175	0.95	1.118
2	-0.63	2.72	1.123	2.040	1.123	2.040	1.582	0.90	1.431
3	0.14	0.98	1.208	2.595	1.208	2.595	1.902	0.86	1.637
4	-0.09	-0.02	1.240	2.705	1.240	0.232	0.736	0.82	0.603
5	-0.49	0.39	1.175	3.061	1.175	0.000	0.588	0.78	0.458
									5.247

This value is likely to be very inaccurate given the very small number of simulated paths. Using 10,000 trials instead of two gives a more reliable estimate of \$4.6 million. This is more than the estimate for the annually capped scheme in Example 5 (\$4.5 million), which illustrates that a T -year scheme with a cap of $\$D$ in total is usually more expensive than a scheme with a cap of $\$(D/T)$ per year. One reason for this is that a total expenditure cap only starts to become material in the later years of the scheme when the dollars saved are discounted more heavily and are thus less valuable.

5.5 Access subsidies and multi-year voucher schemes

Access subsidies could be modeled in the same way as connection subsidies, by treating the number of customers with access as the random variable of interest, and assuming it evolved according to geometric brownian motion. This approach may work best when there is significant probability that connections made in one year will be discontinued in later years (perhaps because of a lack of maintenance, perhaps because the customer does not pay).

When access in one year depends chiefly on whether a connection was made previously, it is possible to analyze access as cumulative connections made in previous years. In this case, the rate of new connections can continue to be considered the underlying risky variable and a modified version of the analysis set out earlier is necessary. (This method can also be used when disconnection is not

improbable; in this case the approach described below provides an estimate of an upper bound to the true cost.)

For valuation purposes, the crucial point to note about such schemes is that X_t units of output at date t commits the government to a payment of sX_t for the remainder of the scheme. That is, when the firm produces a unit of output at date t , we can treat it as having initiated a riskless annuity of s dollars for a further $T - t$ years. Hence, the cost of an access scheme can be estimated by calculating the discounted sum of the initiation-date values of a series of annuities (see Example 4).

A slight variation on the above scheme occurs when the duration of the subsidy differs from the remaining life of the scheme. For example, a government may offer to provide students with vouchers that make some fixed dollar contribution to education costs for n years, with these vouchers to be available to new students for the next T years. That is, any student who begins the system of education being subsidized anytime in the next T years subsequently receives the subsidy for n years from the date of entry.

In this case, each new student commits the government to payment of an n -year annuity. Thus, if the number of new students in the first year is X_t and the dollars per student is s , then the government is effectively committed to payment of an n -year riskless annuity of sX_t beginning in year t . As before, the cost of the scheme can be estimated by calculating the discounted sum of a set of n -year annuities.

Access and voucher schemes can also be subject to annual or cumulative caps, and these can be valued using Monte Carlo simulation in exactly the same way as illustrated earlier.

6 Some additional valuation issues

6.1 Varying growth rates

Throughout section 5, we assumed that the underlying variables follow a geometric brownian motion, and therefore that the expected rate of growth in these variables is a constant. While this is sometimes a reasonable approximation, it may not always be. For example, the introduction of an output-based aid scheme may be expected to provide an initial spur to output, with this effect leveling off as time goes on. In such a case, the expected growth rate would not be constant, but would instead start at a high level and then fall.

Fortunately, this is not a complication that is, in general, difficult to deal with. As section 4 explained, the certainty-equivalent method used in valuing payment schemes does not require that growth rates be constant. So long as the growth rate for each date can be identified, so can the corresponding certainty-equivalent payment.

6.2 Mean-reverting processes

In general, the same approach can be followed for any stochastic process describing X . For example, suppose X follows a mean-reverting process

$$\frac{dX}{X} = \phi(\bar{X} - X)dt + \sigma dz$$

where \bar{X} is the “normal” level of X (that is, the level towards which X tends to revert) and ϕ is the speed of reversion. In this

case, the time- t growth rate depends on the entering value of X , and thus on past growth rates. It follows that²⁰

$$X_t = X_{t-1} \exp\left\{\left(\left(\varphi(\bar{X} - X_{t-1}) - \frac{\sigma^2}{2}\right) + \sigma Z_t\right)\right\}$$

This equation can be used to calculate certainty-equivalent values for X and Y in the usual way by using Monte Carlo to simulate alternative paths for X using the risk-adjusted growth rates $\varphi(\bar{X} - X) - \lambda$.²¹

6.3 Non-CAPM risk premiums

A more knotty problem concerns estimation of the appropriate risk adjustment to the expected growth rate. Throughout section 5, we assumed that this was given by the CAPM, but empirical evidence from financial markets casts doubt on the accuracy of this assumption.

It is important to remember that the CAPM is a model of financial market prices. Additional assumptions are needed to apply it to nonmarket variables (as we did in section 5), but if we are to use it for the latter purpose then we would like to at least be confident that it is able to perform its primary role of explaining market pricing. Unfortunately, finance researchers have become increasingly skeptical about its ability to do this. This modification of the profession's views is neatly summarized by Cochrane (2000):

We once thought that the CAPM provided a good description of why average returns on some stocks, portfolios, funds, or

²⁰ This is a result of Itô's lemma. See footnote 6.

²¹ In principle, similar comments apply to the case where volatility is not constant. However, the intricacies involved in the modeling and estimation of time-varying volatility are a topic in themselves; see, for example, Campbell and others (1997).

strategies were higher than others. Now we recognize that the average returns of many investment opportunities cannot be explained by the CAPM..." (page 36)

In sum, it now appears that investors can earn a substantial premium for holding dimensions of risk unrelated to market movements..." (page 56)

Faced with these claims, finance researchers have recently spent a great deal of time considering the possibility that the CAPM is deficient, despite an initial, and understandable, reluctance to do so.

One reason concerns the role of nontraded assets. Contrary to the assumptions of the CAPM, many assets are not tradable and this increases the demand for those traded securities that hedge the returns on the non-traded assets. Such securities will have higher prices, and therefore lower expected returns, than is predicted by the CAPM. For example, most investors have a job, but the human capital tied up in this activity plays no role in the CAPM. Instead, matters are simplified by assuming that investors care only about the returns on their investment portfolio. However, in reality, employment returns matter as well, as events like recessions hurt most investors: even if they don't lose their jobs, they tend to get lower salaries or other compensation. As a result, investors should prefer stocks that do well in recessions, since these offset the fall in their employment income. Consequently, procyclical stocks should have to offer higher average returns than countercyclical stocks of the same beta. That is, expected returns should depend on the covariation with recessions as well as with the market return.

More generally, an asset's risk reflects the extent to which its returns do poorly in "bad times"—times in which investors particularly wish their investments to do well. Although the market return is indeed one indicator of "bad times", there are potentially many others, for example, recessions, interest rates, and any

variables that provide information about investment opportunities (such as the price-dividend ratio or the term structure).

A generalization of the CAPM that incorporates some of these ideas is the Consumption Capital Asset Pricing Model (CCAPM) of Breeden (1979). This provides us with an alternative way of estimating the risk premium

$$\lambda = \beta_c [\gamma \sigma_{\Delta c}^2] \quad (13)$$

where

$\beta_c = \frac{\text{cov}(\%Y, \%c)}{\sigma_m^2}$ is the consumption “beta” of the payment stream Y

c is per-capita aggregate consumption

$\sigma_{\Delta c}^2$ is the variance of percentage changes in per-capita aggregate consumption

γ is a measure of the market’s aversion to risk.

Unfortunately, despite its theoretical superiority, the empirical performance of the CCAPM is even worse than that of the CAPM, so using (13) to estimate risk-adjusted growth rates is hard to justify.

The main difference between the CAPM and the CCAPM lies in their representations of the single aggregate factor that determines all pricing: the market portfolio in the CAPM and per-capita aggregate consumption in the CCAPM. Two models that allow for multiple priced factors are the Arbitrage Pricing Theory (APT) and

the Intertemporal Capital Asset Pricing Model (ICAPM).²² In these models

$$\lambda = \sum_{j=1}^m \beta_j \phi_j \quad (14)$$

where ϕ_j is the premium attached to factor j and β_j is the beta of the payment stream with respect to factor j . Unfortunately for practical purposes, neither the APT nor the ICAPM identify the priced factors. In the APT, the factors are simply any variables that bear a systematic relationship to market prices, while the factors in the ICAPM are any variables that bear a systematic relationship to investment opportunities. Thus, practical implementation of (14) is largely a matter of educated guesswork.

The most popular example of such guesswork is that of Fama and French (1993) who suggest a three-factor model that in addition to the CAPM beta also includes size and the book-to-market ratio as additional explanatory variables.

$$\lambda = \beta_1(R_m - r) + \beta_2SMB + \beta_3HML$$

where *SMB* is the size factor (the return on small firms less the return on large firms) and *HML* is the book-to-market factor (the return on high book-to-market firms less the return on low book-to-market firms).

However, despite initial enthusiasm for this model, two problems have become apparent. First, the evidence is mixed on the extent to which the three-factor model performs better than the CAPM in explaining variation in market returns.²³ Second, the interpretation

²² The APT is due to Ross (1976) and the ICAPM to Merton (1973). An introduction to the former can be found in Brealey and Myers (2000); Cochrane (2001) provides an excellent discussion of the latter.

²³ See Bartholdy and Peare (2003).

of the additional factors is problematical. For example, a high book-to-market ratio could indicate an additional risk factor (“financial distress”), as claimed by Fama and French. If so, the three-factor model does indeed capture a component of an investment’s risk that is missing from the CAPM. As Stein (1996) notes, the inevitable conclusion from this interpretation is that “one must throw out the CAPM and in its place use the (three-factor) model to set hurdle rates.”

Alternatively, however, a high book-to-market ratio could simply indicate market mispricing. That is, an asset with a high book-to-market ratio is one that is under-valued relative to fundamentals. Over time, the undervaluation disappears and in this sense the high book-to-market ratio “predicts” the high realized returns.

In these circumstances it’s hard to have much confidence that the three-factor model will provide an accurate measure of the true cost of risk. The standard argument for doing so requires that there be a one-to-one link between the market’s expected return on an asset and the fundamental risk of that asset. When this link no longer holds (that is, when there is mispricing), it’s not clear that the best estimate of expected return should be used to set discount rates. For example, if a high book-to-market ratio reflects not high risk, but undervaluation, it makes little sense to set a high discount rate as implied by the three-factor model.

In these circumstances, Stein (1996) shows that it is generally better to use the CAPM (even though it offers an inferior estimate of expected future returns) if decision makers are primarily interested in the long-run economic consequences of their actions. The intuition is that the CAPM continues to provide a reasonable measure of an asset’s long-run risk, even in circumstances where it is unable to predict short-term returns.

To sum up, the CAPM is undoubtedly flawed. However, more-complex models appear to have largely similar problems and thus

do not seem to justify their additional cost. Despite its problems, the CAPM is, given our current state of knowledge, still the main game in town.

References

- Bartholdy, J and P Peare, 2003. Estimation of expected return: CAPM versus Fama and French. Aarhus Business School working paper.
- Baumol, William J. 1963. An expected gain-confidence limit criterion for portfolio selection. *Management Science* 10, 174-182.
- Boyle, Phelim. 1977. Options: a Monte Carlo approach. *Journal of Financial Economics* 4, 323-338.
- Brealey, Richard A. and Stewart C. Myers. 2000. *Principles of Corporate Finance* (6th Edition). International: McGraw-Hill.
- Breeden, Douglas T. 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265-296.
- Brook, Penelope J., and Murray Petrie. 2001. Output-base aid: Precedents, promises, and challenges. In Brook and Smith 2001.
- Brook, Penelope J., and Suzanne M. Smith. 2001. *Contracting for Public Services: Output-Based Aid and Its Applications*. Washington, D. C.: World Bank.
- Campbell, John Y, Andrew W. Lo, and Craig Mackinlay. 1997. *The Econometrics of Financial Markets*. Princeton: Princeton University Press.
- Cannock, Geoffrey. 2001. Expanding rural telephony. In Brook and Smith 2001.
- Cochrane, John. 2000, New facts in finance, Federal Reserve Bank of Chicago Economic Perspectives Economic Perspectives, 36-57
- _____, 2001. *Asset Pricing*. Princeton: Princeton University Press.

- Cox, John C., Jonathan E. Ingersoll, and Stephen A Ross, 1985. An intertemporal model of general equilibrium asset prices. *Econometrica* 53, 363-384.
- Cox, John C. and Stephen A Ross. 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145-166.
- Dowd, Kevin. 1998. *Beyond Value At Risk: The New Science of Risk Management*. Chichester: Wiley.
- Drees, Franz, Jordan Schwartz, and Alexander Bakalion. 2004. Output-Based Aid in Water: Lessons in Implementation from a Pilot in Paraguay. *Public Policy for the Private Sector*. 270.
- Duffie, Darrell. 1996. *Dynamic Asset Pricing Theory*. Princeton: Princeton University Press.
- Fama, Eugene, and Kenneth French, 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33, 3-56.
- Gómez-Lobo, Andrés. 2001. "Output-Based Consumption Subsidies in Chile". Chapter 2 of Brook and Smith 2001. Available at <http://www.gpoba.org/>.
- Harris, Clive. 2002. Network Expansion Using an Output-Based Scheme in Guatemala. *Public Policy for the Private Sector*. 245. Available at <http://www.gpoba.org/>.
- Hull, John C. 2003. *Options, Futures, and Other Derivatives* (5th ed.). New Jersey: Prentice-Hall.
- Liautaud, Gerard. 2001. Maintaining roads: Experience with output-based contracts in Argentina. In Brook and Smith 2001.

- Lintner, J. 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47, 13-37.
- Marin, Philippe. 2002. Output-Based Aid: Possible Applications in the Design of Water Concessions. World Bank, Private Sector Advisory Services. Available at <http://www.gpoba.org/>.
- Merton, Robert C. 1973. An intertemporal capital asset pricing model. *Econometrica* 41, 867-887.
- Olsen, Robert A. 1997. Investment risk: The experts' perspective. *Financial Analysts Journal* 53, 62-66.
- Ritchken, Peter. 1987. *Options*. Illinois: Scott, Foresman & Co.
- Ross, Stephen. 1976. The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13, 341-60.
- Scott, Graham, and Douglas Birnie. 2002. Indicative results and lessons learned from output-based funding trials of public transport services in New Zealand. Available at <http://www.gpoba.org/>.
- Sharpe, William. 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425-42.
- Smith, Warrick. 2001. Designing output-based aid schemes: A checklist. In Brook and Smith 2001.
- Stein, Jeremy C. 1996. Rational capital budgeting in an irrational world. *Journal of Business* 69, 429-455.

GPOBA

The Global Partnership on Output-Based Aid

About the GPOBA Working Paper series:

The GPOBA Working Paper Series is a forum for presenting topics related to output-based aid approaches that may help in enhancing the design of future OBA projects. The series will focus on the infrastructure and social services sectors, but may also include papers that are non-sector specific. The forum is

intended to be a relatively informal and efficient way to discuss relevant OBA design issues. The conclusions presented in each working paper are those conclusions of the authors, and do not necessarily represent the views of GPOBA or its donors.

Global Partnership on Output Based Aid

World Bank
Mailstop: H3-300
Washington, DC 20433, USA

To find out more, visit
www.gpoba.org