



REALIZING THE POTENTIAL FOR PROFITABLE INVESTMENT IN AFRICA
High-Level Seminar organized by the IMF Institute and the Joint Africa Institute
TUNIS, TUNISIA, FEBRUARY 28–MARCH 1, 2006

**Public risk in private
infrastructure**

Timothy Irwin
Senior Economist
Infrastructure Economics and Finance
World Bank

Paper presented at the high-level seminar: *Realizing the Potential for Profitable Investment in Africa*
Organized by the IMF Institute and the Joint Africa Institute
Tunis, Tunisia, February 28 – March 1, 2006

The views expressed in this paper are those of the author(s) only, and the presence of them, or of links to them, on the IMF website does not imply that the IMF, its Executive Board, or its management endorses or shares the views expressed in the paper.

Draft

Public risk in private infrastructure

Timothy Irwin

Draft

December 2005. Please send comments to tirwin@worldbank.org.

Contents

Preface 5

1 Overview 7

1.1 Allocation 11

1.2 Valuation 14

1.3 Rules 15

2 Two centuries of guarantees 19

2.1 Guarantees for bridges and canals 19

2.2 French railway guarantees 21

2.3 Railway guarantees in the rest of the world 26

2.4 Kilometric guarantees in Turkey 28

2.5 Debt guarantees and long-term purchase contracts 32

2.6 Modern revenue guarantees 35

2.7 Exchange rates, construction costs, and rates of return 37

3 Obstacles to good decisions 41

3.1 Cognitive obstacles 41

3.2 Political obstacles 54

3.3 Overcoming the obstacles 57

4 Allocating exposure to risk 59

4.1 Definitions 59

4.2 A principle 71

4.3 Governments' special features 82

5 The allocation of three risks 87

5.1 Exchange-rate risk 87

5.2 Insolvency risk 95

5.3 Policy risk 105

6 Rules 123

6.1 Involving people with an interest in future costs 125

6.2 Charging fees 127

6.3	Setting criteria and requiring analysis	129
6.4	Using markets to value guarantees	133
6.5	Reporting according to modern accounting standards	135
6.6	Disclosing additional information	143
6.7	Budgeting for guarantees	146
6.8	Managing cash and the chance of extreme losses	154
7	Valuing exposure to risk	157
7.1	Identifying exposure	158
7.2	Measuring exposure	160
7.3	Valuing exposure	173
8	The valuation of exposure to three risks	189
8.1	Exchange-rate risk	189
8.2	Insolvency risk	195
8.3	Policy risk	209
	Appendix to Chapter 3	217
	Appendix to Chapter 7	221
	References	231

Preface

Many governments want private firms to finance new infrastructure. Before the firms agree, however, they usually ask the government to bear some of the risks. They might, for example, ask the government to compensate them if demand falls short of forecasts or to promise to repay their debts if they become insolvent. At the very least, they might want the government to allow them to charge a certain price or else compensate them.

This book aims to help governments respond to such requests. As well as seeking to make precise the oft-invoked principle that risks should be allocated to those best placed to manage them, it explains how governments can value the guarantees they are thinking of granting and how they can modify aspects of public-sector management to improve the likely quality of their decisions about guarantees.

Governments need to have an integrated policy toward guarantees, so the book tries to take an integrated approach to the subject—discussing allocation, valuation, and management in the same terms. It also draws on a diverse range of disciplines including finance, history, accounting, economics, and psychology. This imposes an inevitable cost in specialization, but the alternative seems worse, since no single discipline can answer the questions facing governments.

Although intended mainly for governments and those who advise them, the book may be of interest to others, since the problems of allocating and valuing exposure to risk are not specific to governments. For similar reasons, although its focus is physical infrastructure, it may be of interest to people working on public-private partnerships in education, health, and other social services.

1 Overview

Governments sometimes ask private firms to finance new infrastructure projects and, in return, agree to bear some of the risks. The approach is appealing, because the government need pay nothing immediately and can benefit from the skill and enterprise of private firms. But it can cause problems.

In the 1990s, for example, the Korean government guaranteed a privately financed road linking Seoul to a new airport at Incheon 90 percent of a twenty-year forecast of revenue. The government didn't have to pay anything upfront and would get to keep any revenue exceeding 110 percent of the forecast. When the road opened in 2000, however, traffic revenue turned out to be less than half the forecast. As a result, the government has had to pay tens of millions of dollars every year. How much it will have to pay over the life of the guarantee is uncertain; as a present value, it may be as much as \$1.5 billion.¹

The government's policy may not have been wrong, but it does raise questions. Should the government really have borne demand risk in the project? Could it have estimated the cost of its guarantee before granting it? If so, should it have disclosed an estimate of the cost in its accounts? More generally, could the government have built the road more cheaply using public finance? Or would it have been better to use private finance without a revenue guarantee, if necessary giving the firm a straightforward subsidy.

These questions are hard to answer even though government guarantees² have been around for a couple of hundred years. Most

¹ Irwin (2004).

² We use "guarantee" to mean "something that assures a particular outcome or condition," (*American Heritage Dictionary*, fourth edition). A "guaranty," on the

notably, government guarantees were used in the nineteenth century to encourage the private financing of railways. Argentina, for example, guaranteed railway investors returns of six or seven percent on the capital they invested. (On one account, the investors who got the more attractive guarantee did so “not because they necessarily had a better case, but because they were more generous in their dishing out of bribes.”³) The guarantees helped Argentina attract investment from foreign capital markets and reflected a view that Argentina had to compete for such funds by offering incentives like those offered by other countries.⁴ Yet the government didn’t always have enough money to meet its commitments, in part because of the difficulty of accurately budgeting for claims and in part because the government usually had to make larger payments just when its tax revenue was low.⁵ In time, the guarantees contributed to a fiscal crisis:

As the principal official obligation, railway guarantees were ... largely responsible for the crisis of confidence which brought the whole fragile edifice of the Argentine economic miracle to ruin ... Many of the companies formed during this period were committed to long-term dependence upon the guarantee. Indeed ... it might be argued that many were projected merely for the purpose of obtaining a guarantee.⁶

Guarantees have sometimes been used with better results. The Chilean government has given many revenue guarantees and a few exchange-rate guarantees to privately financed toll roads. The revenue guarantees typically ensure the concessionaire gets revenue equal to 70 percent of the estimated present value of its costs, including the costs of investment, operations, and maintenance; the guaranteed revenue might be spread over twenty years, providing as much as 85 percent of forecast revenue in early

other hand, is “an agreement by which one person assumes the responsibility of assuring payment or fulfillment of another’s debts or obligations.”

³ Anthony Burton (1994, page 122).

⁴ Colin Lewis (1983, page 11).

⁵ Lewis (1983, pages 102, 105, and 122).

⁶ Lewis (1983, page 86).

years and less later on. So far, the government has got a great deal of investment without having to pay much because of these guarantees.⁷

Even in Chile, however, the guarantees raise difficult questions. A recession would cause traffic to grow more slowly than expected, possibly triggering many guarantees just when tax revenue was weak. One might therefore ask, how serious is the risk, and how can the government measure it? What is the value of its outstanding liability? And should it be planning now for the possibility of future payments?

The questions raised by guarantees are perhaps most pressing in developing countries, but they are not unique to them. To take just one example, the state of New South Wales in Australia gave a revenue guarantee in the early 1990s to the Sydney Harbour Tunnel, a project developed as an alternative to the Sydney Harbour Bridge. Toll revenue was expected to be too low to cover the tunnel's costs, but by the terms of an "ensured revenue agreement," the government contracted to give the tunnel company a specified amount less tolls on the tunnel.⁸ Thus the government, not the company, bore demand risk. The Auditor-General of New South Wales concluded that the project was more public than private and that, for accounting purposes, the tunnel and associated liabilities were the government's. It qualified its audit of the financial statements of the government agency promoting the road and argued that the agency had chosen nominally private but effectively public finance partly as a way of circumventing a cap on public borrowing.⁹

It is difficult for governments to make good decisions about guarantees. To start with, there is no agreement among experts and

⁷ Andrés Gómez-Lobo and Sergio Hinojosa (2000) describe the guarantees and their results until 2000.

⁸ Government of New South Wales, Australia, Auditor-General's Office (1994, page 279).

⁹ Government of New South Wales, Australia, Auditor-General's Office (1994).

advisers on which risks governments should bear in privately financed projects. Should they bear demand risk in toll road projects? Should they give exchange-rate guarantees when investors borrow in foreign currencies? Or should they shield investors from exchange-rate risk by increasing the price of the service when the local currency depreciates? Should they protect creditors from losses in the event of the project's being terminated? Should they compensate investors for changes in government policy? All changes? Some changes but not others? No changes? Advisers have opinions on these subjects, but there is no consensus.

The difficulty of knowing how best to allocate risks is compounded by other problems. First, politics can encourage governments to bear more risk than is in the public interest. Governments are buffeted on all sides by proposals for subsidies, but unless the beneficiaries are widely regarded as deserving, the most transparent such proposals tend to fail. Successful proposals tend to have opaque costs and to come with a rationale explaining how they are good for the country and don't merely redistribute value. Proposals for guarantees can meet these criteria, especially when the government's accounting and budgeting fail to recognize their costs. They come with plausible rationales about risk sharing, and taxpayers are unlikely to understand the costs.

Second, government decisions would be difficult even in the absence of political pressures. Psychological research shows that people struggle to make accurate judgments about risks and then fail to make the best use of even their imperfect judgments. Most people, for example, are overconfident in their judgments, and therefore think the world is more predictable than it is. Government decision makers may fall into the same trap, underestimating the risks to which they are exposing the public when they issue guarantees. They may also make decisions about guarantees that are irrational given their judgments. Research shows that people can switch from being risk averse to risk seeking, just because the framing of a choice changes. They can also be irrationally risk averse when they consider risks one by one, instead of thinking of their total portfolio of assets and liabilities.

In sum, governments can easily make poor decisions about guarantees. There is no simple solution to this problem, but good decisions are more likely if three conditions are met:

- (i) the government's advisers and decision makers have a framework for judging when a guarantee is likely to improve a project;
- (ii) the government's advisers know how to estimate the cost of a guarantee; and
- (iii) the government's decision makers must follow rules that encourage careful consideration of a guarantee's costs and benefits.

Helping governments fulfill these three conditions is the aim of this book.

1.1 Allocation

To judge when governments should bear risk, we need a framework for deciding the more general question of how risk should be allocated. And to set out such a framework we need to clarify what we mean by risk.

As we will use the term, *risk* is unpredictable variation in value. It includes the possibility of unexpectedly good as well as unexpectedly bad outcomes. The *risk of a project* is unpredictable variation in the total value of the project, taking account of not only the value of the project company but also the value accruing to customers, the government, and other stakeholders. A *stakeholder's risk* in a project is unpredictable variation in the value of the stakeholder's interest in the project.

Particular risks can also be defined. *Demand risk* is unpredictable variation in value arising from unpredictable variation in demand. *Construction-cost risk* is unpredictable variation in value arising from unpredictable variation in construction costs. Thus we use "risk" to refer both to the total risk of the project and to the

components of such risk. We describe the sources of risk as *risk factors*.

How should risks be allocated? The conventional answer is that each risk should be allocated to the party best able to manage it. The answer looks right, but it is too vague to be very helpful by itself. The following principle tries to clarify it: Each risk should be allocated, along with rights to make related decisions, so as to maximize project value, taking account of each party's ability to

- (i) influence the corresponding risk factor,
- (ii) influence the sensitivity of total project value to the corresponding risk factor, for example by anticipating or responding to the risk factor, and
- (iii) absorb the risk.

The principle refers to three ways in which a risk can be managed. First, there are times when someone can influence the risk factor—that is, can take action to improve or worsen the risky outcome. For example, a construction company can change construction costs by its choice of materials and techniques and the way it pays and manages its workers. If no one else can manage construction-cost risk, the principle says that the risk should be allocated to the construction company. Such an allocation doesn't eliminate the risk; the construction company's profits remain uncertain. But, compared with other allocations, it will tend to lower the cost of construction and increase the total value of the project.

Second, there are times when someone can influence the sensitivity of the value of the project to the risk factor. For example, no one can influence whether an earthquake occurs, but by carefully choosing the site of a project, the firm or the government may be able to reduce the loss that an earthquake would cause. If the risk is thus anticipated, the value of the project is made less sensitive to the risk factor. Someone may also be able to respond to variation in a risk factor. A firm may be able to switch between inputs as their relative price changes, mitigating downside risk in the price of one input and exploiting upside risk in the price of the other. The

principle says that, other things equal, the party that can best anticipate or respond to the risk factor should bear the risk

Third, there are times when no one can influence, anticipate, or respond to the source of a risk in a way that changes the project's value. At such times, the risk should be allocated to the party that can absorb the risk most easily, or in other words, bear the immutable risk at the lowest cost. Customers, for example, may be able to absorb the risk of inflation in the price of a service because their incomes are adjusted for inflation. The firm or the government may be able to absorb a risk because it can buy derivatives or insurance to protect it from the risk. And the shareholders who ultimately bear risk allocated to the firm may be able to absorb a risk simply because they have well-diversified portfolios.

Applying the principle of risk allocation to a particular government's decision to bear a particular risk in a particular project can be hard. The details of the government, the risk, and the project matter. Trying to give definitive general advice on whether governments should bear particular risks is thus futile. Without trying to be definitive, however, subsequent chapters argue that governments should be inclined to bear project-specific risks that they control or strongly influence, such as risks related to prices and quality standards that they regulate. One way governments can do this is to contract with the firm to set prices or quality standards in a particular way and therefore to assume an obligation to compensate if they change their mind. Governments might sometimes strongly influence other risk factors, such as demand for a road when demand depends heavily on the construction of competing and complementary roads in a government-planned network. In such a case, it might make sense for the government to bear demand risk, by giving a revenue guarantee, financing the road itself, or promising payments independent of demand to the owner of an untolled road.

The subsequent chapters also argue that governments should be disinclined to bear exchange-rate and other economy-wide risks. Although governments can often influence such risks, they shouldn't usually shape economy-wide policy to suit the interests

of a particular project. Moreover, although the firm and its creditors cannot influence economy-wide risk factors, they can often influence the sensitivity of the project's value to the risk factor. Their choice of the extent of borrowing in foreign currency, for example, influences how sensitive the value of the project is to the exchange rate. For similar reasons, governments shouldn't usually bear the risk of the firm's insolvency—that is, shouldn't usually issue debt guarantees.

The principle also says that the risks a government should bear depend on the way it allocates rights to make related decisions. The more rights it devolves, the more risk it can reasonably transfer. Conversely, which rights a government should retain depends on the risks it chooses to bear. If a government offers an exchange-rate guarantee, for example, it should not give the firm the right to choose freely its borrowing in foreign currency.

1.2 Valuation

An understanding of when guarantees can improve a project won't ensure good decisions. Politics and psychology will still get in the way. To buttress the benefits of understanding, governments can estimate the cost of the guarantees they are thinking of giving—or, to use a different expression, can value those guarantees. Valuation allows policy-makers to supplement their own susceptible intuition with quantitative estimates, reducing the temptation to bear risk when it is unlikely to increase total project value.

Although some of the details of valuation are technical, the main ideas are simple. The first step is to identify the risks the government is thinking of bearing. What are the rights and obligations, that is, that will cause the government to receive or make payments? And, given those rights and obligations, what are the risk factors (demand, construction costs, and so on) that will determine what the government pays or receives?

Given a model of those risk factors, the next step is to measure the government's exposure to risk; that is, to answer questions such as these: What is the most the government might lose? How much can it expect to lose? And what is the chance of its making certain large

losses? To make this concrete, suppose the government tosses a coin four times and offers to pay the firm \$1 every time the coin lands head up. The most it could lose is \$4. The amount it can expect to lose—what it would lose on average if it played the game many times—is \$2. And the probability of its losing, say, \$4 is one in sixteen ($\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2}$).

The third step is to estimate the cost of bearing the risk. To do that, the expected payment needs to be adjusted to take account of time and risk. Adjusting for the timing of payments is the easier of the two tasks. A payment made in the future is less costly than a payment made now, so the future payment needs to be discounted at an interest rate reflecting the time value of money (the riskless rate of interest).

Adjusting for the risk of payments is harder, and dodging the problem is sometimes justified in practice. But most guarantees are worth more than their expected value discounted at the riskless rate of interest. So a government that ignores the cost of bearing risk will issue too many guarantees. Fortunately, there are methods widely used in financial markets that can often generate a reasonable estimate of the cost of bearing risk.

1.3 Rules

Improvements in decisions about guarantees can be sought case by case; with better knowledge of the principles that should govern the allocation of risk and the ability to value exposure, governments can try to make better decisions in each case they encounter. To fully offset the problems of politics and psychology, however, improvements can also be sought in the rules that govern case-by-case decisions on whether to bear risk. In particular, governments can change those rules to try to ensure that decision makers have access to relevant information and have incentives to act in the public interest.

Public-sector accounting standards are crucial. Cash accounting encourages governments to ignore the costs of decisions that have no immediate effects on their cash expenditures or revenues. Modern accrual accounting standards, on the other hand, require

the recognition of at least some noncash costs. In particular, the best standards require recognition of the liabilities created by some guarantees. And when they don't require the recognition of a liability, they usually require the disclosure of information in notes to the accounts. Two sets of modern accrual standards that governments can adopt are International Public Sector Accounting Standards and the International Monetary Fund's *Government Finance Statistics Manual 2001*.

Budget rules are also crucial. Good budget rules require governments to examine spending proposals simultaneously and therefore confront the tradeoffs inherent in them. Decisions about exposure to risk should be made in the same way. Budgets should give approval to incur noncash costs, as well as to disburse cash, and a dollar's cost incurred by guarantee should count as a dollar spent in cash. Budgeting with good accounting helps do this; but even the best accounting standards disregard the cost of some forms of exposure to risk, and most governments are still some way from reporting to the best standards.

Thus stopgap measures that force the counting of guarantees may help. Some governments have used special funds to make up for weak accounting and to help manage the cash-flow risks of guarantees. When a ministry issues a guarantee, the government can require it to contribute to the fund the estimated cost of the guarantee. The fund can then be used to meet, or contribute toward, payments if the guarantee is called. Another option is to set up a purely paper transaction, in which money goes to the ministry issuing the guarantee and then back to the government's main account. The government continues to manage cash as before, but an estimate of the cost of guarantees shows up in the issuing ministry's budget.

Governments can also change rules to enlist the help of outsiders. Laws requiring the disclosure of contracts and other documents—routinely or after requests made under freedom-of-information laws—give members of the public the opportunity to comment on and criticize the government's decision making about risks. Irritating as it may be for governments, the possibility of external criticism probably improves the average decision. Governments

can also require that some guarantees be issued only for a price—a device that, among other things, means that the recipient of the guarantee is part of the group deciding whether the guarantee will be issued, which may reduce the chance of a guarantee's being granted when its costs exceed its benefits.

These ideas are developed in the chapters that follow. Chapter 4 sets out a framework for allocating risks and therefore for deciding when to grant guarantees, using demand risk as an example. Chapter 5 applies the framework to three other risks: exchange-rate risk, insolvency risk, and policy risk. Chapter 6 develops the argument for changing higher-level rules to improve case-by-case decisions. Then Chapter 7 sets out a framework for measuring and valuing exposure to risks, including those created by guarantees, using demand risk as an example. Chapter 8 applies the framework to exchange-rate risk, insolvency risk, and policy risk. We start by reviewing the history of government guarantees and the cognitive and political obstacles that stand in the way of good decisions.