FONDATION POUR LES ÉTUDES ET RECHERCHES SUR LE DÉVELOPPEMENT INTERNATIONAL

Working Paper 55
November 2012
Development Policies

Savings from Natural Resource Revenues in Developing Countries : Principles and Policy Rules

Paul Collier

Paul Collier is Professor of Economics and Public Policy at the Blavatnik School of Government and the Director of the Centre for the Study of African Economies at Oxford University. His mains areas of activity are governance in low-income countries, especially the political economy of democracy, economic growth in Africa, economics of civil war, aid, globalisation and poverty. He recently published "The Plundered Planet: How to reconcile prosperity with nature" (2010), "Wars, Guns and Votes: Democracy in Dangerous Places" (2009) and "The Bottom Billion" (2008), Estoril Global Issues Distinguished Book Prize. He is Senior Fellow at Ferdi.

Abstract

Many poor countries are now discovering valuable non-renewable natural resources. Unlike most other sources of tax revenue, the government revenues from the depletion of these resources are both unsustainable and volatile. Each of these features implies that the savings rate appropriate for resource revenues should differ from that on other revenues. Further, a discovery is 'news', requiring a transition from a situation which has suddenly become sub-optimal. Such transitions must be expected to generate costs which will themselves affect the optimal savings rate. While the features themselves have long been well-understood, the implications for optimal savings behaviour are surprisingly underdeveloped. A fortiori, the implications for the rules which might be the practical embodiment of these analytic underpinnings are also underdeveloped.

… /…

LA FERDI EST UNE FONDATION RECONNUE D'UTILITÉ PUBLIQUE.
ELLE MET EN ŒUVRE AVEC L'IDDRI L'INITIATIVE POUR LE DÉVELOPPEMENT ET LA GOUVERNANCE MONDIALE (IDGM).
ELLE COORDONNE LE LABEX IDGM+ QUI L'ASSOCIE AU CERDI ET À L'IDDRI.
CETTE PUBLICATION A BÉNÉFICIÉ D'UNE AIDE DE L'ÉTAT FRANCAIS GÉRÉE PAR L'ANR AU TITRE DU PROGRAMME «INVESTISSEMENTS D'AVENIR»
PORTANT LA RÉFÉRENCE «ANR-10-LABX-14-01».

… / … Our analysis proceeds in four stages. In Section 2 we discuss how savings should optimally respond to the known depletion path of a finite natural resource. In Section 3 we discuss revenue volatility in the context of costly adjustments in spending. In Section 4 we discuss the implications of costs of transition. Finally, in Section 5 we discuss the implications for fiscal rules and institutions. We argue that where revenues from depleting resources are significant, fiscal optimization requires that the intra-temporal rules conventionally incorporated into annual budgeting be supplemented by inter-temporal rules. Such rules protect the efficient fiscal allocation which in conventional circumstances integrated budgets are designed to achieve.

1. Saving in Response to Depletion[1]

In this Section we first consider the response to depletion, deriving optimal savings rules from a family of models based on the concept of permanent income. We progressively incorporate salient complexities which are omitted in the basic permanent income framework.

The Base Case

The extraction of mineral resources is unsustainable because endowments are not renewable. Hence, unless some of the revenues are used to acquire revenue-generating assets, the consumption financed from resource depletion is also unsustainable. The permanent income model seeks to solve for that savings behaviour out of resource revenues that achieves the highest sustainable level of consumption.

We will treat the path of resource extraction as exogenous to government policy: resource extraction companies determine the path of the volume of resources to be extracted and international markets determine the path of prices. While the government is free *ex ante* not to permit extraction, the question as to the appropriate savings rate only arises because the government has agreed to permit it.

To establish the permanent income afforded by a temporary stream of natural resource revenue we first introduce some notation. Let natural resource depletion create a flow of income x_τ for a period to time $\tau \in [0, T]$, after which the resource is exhausted. The present value of the flow remaining at date t is $X_t \equiv \int_t^\infty x_\tau e^{-r_X(\tau-t)} d\tau$, where is r_X the rate at which expected future

[1] For the figures and the mathematics in this section I am indebted to Tony Venables.

resource revenue is discounted. The present value of the windfall at date 0 is X_0, and the present value changes through time according to $\dot{X} = r_X X - x$.

Some of the resource flow is used to accumulate a non-resource asset, A. Hence, total wealth due to the discovery is $W = A + X$. Non-resource assets are accumulated according to :

$$\dot{A} = x + r_A A - c \,. \tag{1}$$

The first term on the right-hand side is the flow of revenue, the second is income on accumulated non-resource assets yielding return r_A, and the third is the consumption financed by resource revenues, c. Thus, adding the two elements of wealth, its evolution is given by :

$$\dot{W} = \dot{A} + \dot{X} = [r_A A + x - c] + r_X X - x = [r_A W - c] + X[r_X - r_A]\,. \tag{2}$$

Our first result is to show that the appropriate savings rate from such natural resource revenues depends upon the horizon to depletion. *The shorter the horizon until expected depletion the higher should be the savings rate*. This follows straightforwardly from the permanent income framework. For given annual revenue, the shorter is the extraction period the lower is the present value of the resource endowment and hence permanent income – the sustainable increase in consumption. With less consumption warranted from a given resource revenue, the higher is the savings rate. An important corollary is that for given constant annual revenue, *the savings rate should rise as resources are depleted*. Each year the horizon to full depletion is shorter and so the appropriate savings rate is higher.

More formally, denote the time preference rate by ρ. For the PIH to apply, economic actors optimize by setting this rate equal to both the return on assets and the rate at which future revenues are discounted, $\rho = r_A = r_X$ ($\equiv r$). Given this equality, sustainable consumption is thus maximized by setting $c = rW$. Using this in equation (2) above, wealth is constant through time at value $W = X_0$. Hence, the (constant) path of consumption, c, is proportional to the initial present value of the windfall:

$$c = rX_0. \tag{3}$$

Since consumption is completely independent of the time profile of revenues, the path of asset accumulation, A, fully accommodates any deviation of the actual revenue profile from its permanent value. The path of savings follows immediately. Saving is $\dot{A}$, and we define s_X as the proportion of resource revenue that is saved at each instant of time, $\dot{A} = s_X x$. Since $\dot{W} = \dot{A} + \dot{X} = s_X x + rX - x = 0$, it follows that the savings rate along the permanent income path is:

$$s_X = 1 - rX / x. \quad (4)$$

Equation (4) formally states our results that for a constant path of revenues, the savings rate should rise as resources are depleted, and should be higher the smaller the discovery: we term this *the depletion effect*. A simple way of thinking about the depletion effect is that as depletion occurs, the stock of non-resource assets gradually accumulates. In turn, this rising stock generates a rising sustainable income, all of which should therefore be used to finance consumption. As a result, the post-discovery level of consumption warranted by PIH can be increasingly financed from these revenues and so less of the (constant) revenues from resource depletion needs to be used for consumption: hence, the rising savings rate.

Figure 1 illustrates two cases with the same overall size of resource discovery. For both the annual revenue from extraction remains constant over the life of the resource, but the revenue levels differ: one has high annual revenue but a short life, the other has lower revenue and a longer life. The left hand panel of Figure 1 shows the two revenue flows, with revenue, *x*, on the left hand panel and time on the vertical. Both profiles, shown by the solid and dashed lines, have the same initial present value, $X_0 = 1$, but that depicted by the solid line depletes the resource slowly so that it lasts for forty years, whereas that depicted by dashed line depletes it more rapidly so that it only lasts for thirteen years. The paths of the savings rate out of resource revenues that are optimal for these depletion profiles are illustrated on the right panel. The two savings rules noted about follow directly from equation (4). First, in both cases the savings rate rises over the course of depletion, since the present value, X, is gradually reducing while the flow *x* and the discount rate r are constant. Indeed, since X goes to zero as exhaustion approaches, the savings rate eventually rises to 100 percent. The second rule is that the savings rate is lower and increasing less fast the longer is the duration of the resource (solid line).

Figure 1: Savings rates from the PIH: Constant Depletion Rates

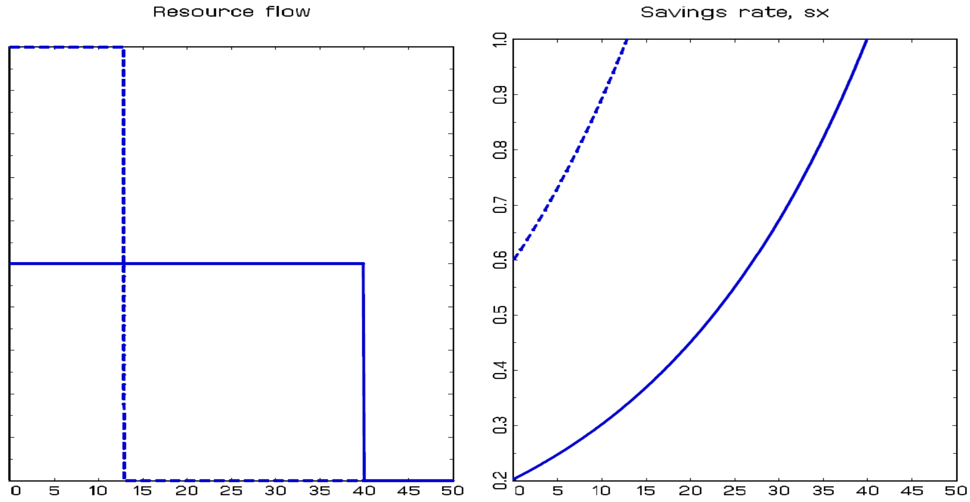

As Figure 1 shows, the depletion effect is dramatic. In the case of slow depletion the optimal savings rate rises from 20 percent in the first year to 100 percent in the final year. Hence, a policy rule which set a *constant* savings rate would be seriously sub-optimal regardless of the level at which it was set. Similarly, the savings path for rapid depletion is strikingly higher than for slow depletion, starting at 60 percent rather than 20 percent. Hence, a policy rule which purported to set the savings rate without reference to the expected time to depletion could not avoid being seriously sub-optimal.

In practice, the most efficient path of extraction is seldom likely to yield constant annual revenue. If the extraction path is changing this generates a further effect on optimal savings which we term *the extraction path effect*. The depletion effect comes from the gradual decrease in the numerator of rX/x to zero at the point of exhaustion. The extraction effect comes from the denominator and can in principle be either positive or negative. If x is rising then the extraction path effect also generates a rising optimal savings rate and so reinforces the depletion effect. If x is declining, then the two effects are qualitatively offsetting.

The more likely case is the former: x is rising over time and so this reinforces the rise in the savings rate. To see this it is helpful to decompose the extraction path into its component parts: the path of the quantity of resources extracted, and the path of their price. While both are, by assumption, exogenous to the government, they are endogenous to the profit calculus of resource extraction companies. The Hotelling Rule predicts that profit maximization by resource extraction companies will generate a path of resource prices which rises by the world interest rate. As we discuss in Section 2.4, the assumptions necessary for this result are implausible. However, even if the Rule provides an incomplete account of the path of resource prices, the behavior on which it rests is undeniably a component, and so it provides some reason to expect that prices will tend to rise. Thus, unless the quantities extracted systematically decline sufficiently rapidly fully to offset this effect, the path of revenues is also likely to be increasing.

Further, even if the quantities extracted decline so rapidly as to more than offset the rise in the price, so that x is declining, the depletion effect will eventually predominate over the extraction path effect. This is because as exhaustion approaches, X tends to zero, and so the optimal savings rate rises to 100 percent. Further, even in the somewhat extreme case in which revenues are declining exponentially, the savings rate does not actually fall, merely remaining constant for a while before rising to 100 percent. This case is illustrated by the dashed lines in Figure 2, the left-hand panel depicting the exponential decline in revenues, and the right-hand panel the path of the savings rate.

The most likely path of quantities extracted from a discovery is hump-shaped, as a phase in which the investment in extraction is building up is followed by one in which extraction becomes progressively more difficult as the most accessible resources have already been removed. If this quantity path eventually overrides the trend for prices to increase, then the extraction path will itself be hump-shaped. This is the case illustrated by the solid lines in Figure 2. As that Figure shows, this actually reinforces the tendency for the optimal savings rate to rise with depletion. Superficially, this is surprising because, whereas during the phase of rising revenues the extraction path effect reinforces the depletion effect, during the phase of declining revenues it correspondingly offsets the depletion effect. However, the reason why the net effect of a humped extraction path is to reinforce the depletion effect is that as exhaustion approaches the depletion effect becomes relatively more important: recall that savings rate from the last revenue generated by extraction should be 100 percent regardless of other considerations. Thus, the most significant impact of a humped extraction path compared with a flat one is its reinforcing effect during the phase in which revenue is rising.

Figure 2: Savings Rules from the PIH: Declining versus Humped Depletion Rates

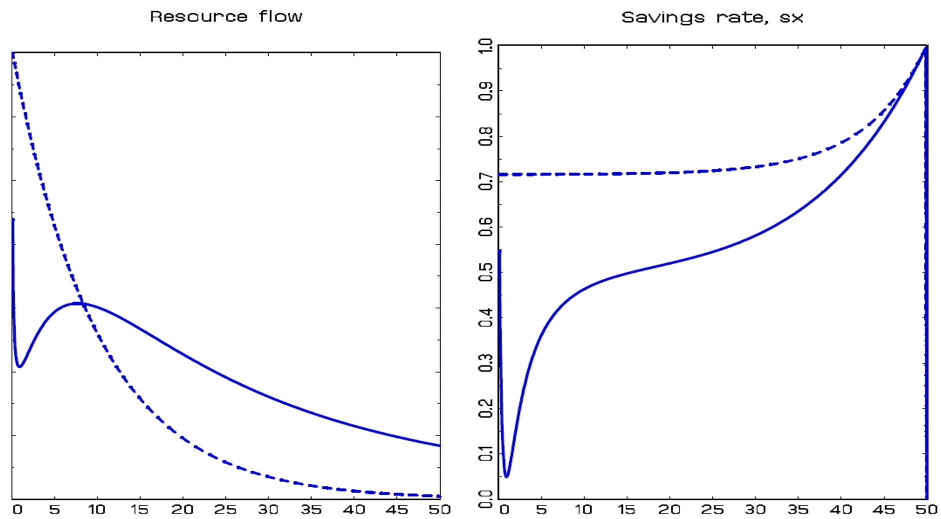

In summary, the optimal savings rate depends upon the extraction path. Two generalizations for policy appear warranted from the above discussion. First, and unsurprisingly, for a given NPV of a resource discovery, the shorter its duration the higher should be the savings rate. Second, and much more surprisingly, for a wide range of extraction paths the optimal savings rate should rise over time. This rising savings rate is generated by two distinct effects, the depletion effect, and the extraction path effect. Typically, though not invariably, the two effects will reinforce each other.

Allowing for Lack of Access to International Capital Markets

The above analysis has made assumptions that may not hold in poor, resource-rich countries. We now begin relaxing these assumptions to see what further effects on optimal savings are likely to be important for such countries.

The first modification we introduce is to capital markets. Most poor countries have only asymmetric access to world capital markets: they can save more readily than they can borrow. Typically, as borrowers they are either entirely cut off – often as a requirement of an IMF Program – or, they can only borrow at a substantial premium over the world risk-free interest rate to compensate for perceived high risks of default. For example, the government of Senegal recently floated an international bond at nine percent, at a time when the world risk-free interest rate was under three percent. An implication of this asymmetric segregation of domestic and international capital markets is that the rate of social time preference prevailing in a poor country need not be the same as in rich ones. Those poor countries that have a lower rate of social time preference than rich ones can acquire claims on rich countries, as China has done vis-à-vis America. But those poor countries that have a higher social time preference than rich ones cannot borrow. The low Chinese social discount rate is atypical for poor countries and may be explained by its equally unusual demographics. Most poor countries have youthful populations that are fully aware of consumption levels elsewhere in the world. This is likely to generate impatience and so a high domestic social rate of time preference. In turn, in equilibrium this will imply that the marginal return on domestic assets is also high.

Hence, if the Hotelling Rule holds, so that the price of natural resources rises at the world rate of time preference, there is a substantial gap between this rate and the domestic discount rate at which future resource revenues are discounted. Stated more formally, $\rho = r_A > r_X$: the domestic rate of time preference and rate of return on assets, though equated, are higher than the global rate of price increase of natural resources.

To see the consequences most clearly, suppose that the global rate of time preference is zero but that the domestic rate of time preference is high. For comparability with our baseline case we will assume that revenue is constant along the extraction path. As extraction proceeds, the time profile of the remaining extraction path shortens: each year the most distant year of extraction is 'lost'. However, as extraction proceeds, the final 'lost' year gets progressively closer. The large difference between the domestic and the international social discount rate has an important implication for the valuation of this sequence of 'lost' years. For example, after the first year of a twenty-year discovery, the 'lost' twentieth year of extraction is not very valuable because, having been discounted at a high rate over a twenty year horizon its NPV was only a small fraction of its eventual value. But after the nineteenth year, the 'lost' year is not the twentieth year of extraction but the second, and so the NPV of the loss is much greater. In consequence, the value of resources remaining in the ground declines at an *accelerating amount*.

In turn, this has unambiguous implications for the savings rate. For PIH, total wealth must remain constant along the extraction path, and so since resource wealth is declining by an accelerating amount, other assets must be accumulated by an accelerating amount. Hence, the amount saved must increase each year. By construction, the value of resource extraction is constant, so that the rate of savings must rise each year.

More generally, the wider is the wedge between the rate of time preference and the rate of change in the price of natural resources the more powerful is this effect. The higher is the gap the faster is the rate of acceleration in the decline in the value of resources remaining in the ground because the larger is the difference between the value of temporally distant extraction and temporally proximate extraction. We term this the *divergent discount rates effect*, and, as with the depletion effect and the extraction path effect, it leads to a rising savings rate over time.

Allowing the Return on Investment to Rise with the Resource Discovery

Above we assumed that although the society was unable to borrow from international capital markets, internally there was a capital market equilibrium in which the social discount rate was equated with the return on domestic non-resource assets. It is reasonable to assume that after a sufficiently prolonged period in which economic conditions remain largely unchanged such an equilibrium will emerge. However, even if the economy is initially so characterized, a major resource discovery is likely to disturb the equilibrium. Specifically, a likely consequence of a resource discovery is that the return on investment rises, temporarily raising it above the rate of time preference, so that $\rho < r_A$. This effect becomes clearer once investment is decomposed into its public and private components, with public capital and private capital as complements. The resource discovery signals to private investors that the domestic market will become larger and so induces an increase in private investment. For example, following the upsurge in resource discoveries in Africa over the past decade, there has been an upsurge in private investment. This correspondingly raises the return on public investment.

To deduce the implications for savings out of resource revenues we need an explicit statement of our (hitherto implicit) framework. The social objective is the present value of the utility of consumption, given in equation (6) below. Consumption expenditure is income (now written to include non-resource income, Y, as well as resource revenue x and asset income $r_A A$) net of saving $\dot{A}$:

$$V = \int_0^\infty u(c)e^{-\rho t}dt = \int_0^\infty u\left(Y + x + r_A A - \dot{A}\right)e^{-\rho t}dt \qquad (5)$$

Before the discovery the economy is stationary, with $x = 0$, zero assets or debt, $A_0 = 0$, and consumption equal to income, $Y = c_0$. Resource revenues flow for a limited period of time, t. From period T the resource is exhausted and so the economy again becomes stationary, with $c_t = Y_t + r_A A_T$. The economic problem is to choose the time path of consumption and asset accumulation during the period when resource revenues are flowing, so as to maximize V subject to these constraints.

The solution is characterized by the Euler equation $\dot{u}_C / u_C = \rho - r_A$ As is conventional, we simplify by assuming that $u()$ is isoelastic, $u = c^{1-1/\sigma}\sigma/(\sigma - 1)$, so that the Euler equation is $\dot{c}/c = \sigma(r_A - \rho)$.

Within this framework, were there no increase in the return on investment, so that $r_A = \rho$, the PIH would follow immediately, with consumption jumping to its new stationary level. This base case is illustrated in figure 3(a). The left hand panel gives the initial and new levels of consumption, the flat upper and lower lines. The central panel depicts the paths of the two assets, with the downward sloping curve giving X, which for simplicity we assume is depleted exponentially. The upward sloping curve is the path of invested assets, A. The combined stock of natural and invested assets, W, is constant (which is necessarily the case with PIH). The right hand panel is the saving rate, (reproducing the path shown by the dashed line in the right hand panel of Figure 2).

With this as the base case, we now introduce the effect of the resource discovery on return on investment. As a result of the increased return the initial equilibrium is disturbed: the rate of return on invested assets now exceeds the rate of time preference ($r_A > \rho$). The resource discovery thus both creates the *need* for more savings to finance the improved investment opportunities, and the surplus that enables this saving to be undertaken while still increasing consumption. In consequence, consumption will be on a rising path, starting below the PIH level and ending above it. This is illustrated in figure 3(b), in the left hand panel of which the dashed line is consumption and the upper solid line the PIH case, for comparison.

Comparing the new savings path with the base case, the jump in the return on investment consequent upon the resource discovery makes little difference to the initial rate of saving (in the example it remains at around 70 percent). However, the savings path now rises over time even though we are working with the special case in which resource revenues are declining exponentially so that in the base case the savings rate is constant. Not only does the savings rate rise over time, there is a phase during which it rises beyond 100 percent, implying that even some of the sustainable revenues from invested assets are saved.

The underlying economics is that the resource discovery has plunged the domestic capital market into disequilibrium. Because it is not able to borrow internationally, this cannot be reconciled by increased borrowing from abroad. There are good reasons to save more, so as to seize the new opportunities for high-return investment, but there are also good reasons to consume more now, given that the high return on investment will enable the economy to grow so that future consumption levels will be higher and the marginal utility of consumption correspondingly lower. Eventually this tension is resolved by the new opportunities for investment winning out: the savings rate is higher. But in the first years of the resource discovery consumption wins out: high-return investments can be postponed, whereas if high-utility consumption opportunities are missed they are lost forever. Hence, at the onset of the discovery it is best to prioritize consumption almost as much as if the return on investment had not risen. As consumption increases, so that there are no further opportunities for high-utility consumption, undertaking the high-return investments becomes increasingly attractive and so the savings rate rises. We term this the *capital markets disequilibrium effect*.

As illustrated in Figure 3, the savings rate rises past 100 percent because, even once it reaches 100 percent the economy is still in disequilibrium, with the return on investment exceeding the rate of time preference. The savings rate therefore rises further to accelerate the restoration of equilibrium.

The above scenario, in which the discovery of natural resources increases private investment and thereby increases the return on public investment, describes the typical recent experience of resource-rich low-income countries.

Figure 3: The resource discovery increases the return on investment

Figure 3a: Baseline PIH, $r_A = \rho$, exponentially declining revenues

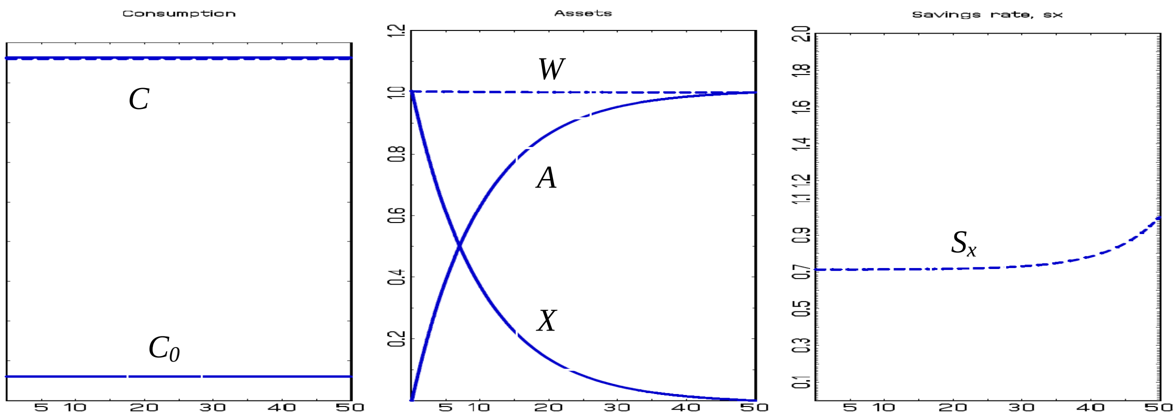

Figure 3b: $r_A > \rho$, exponentially declining revenues

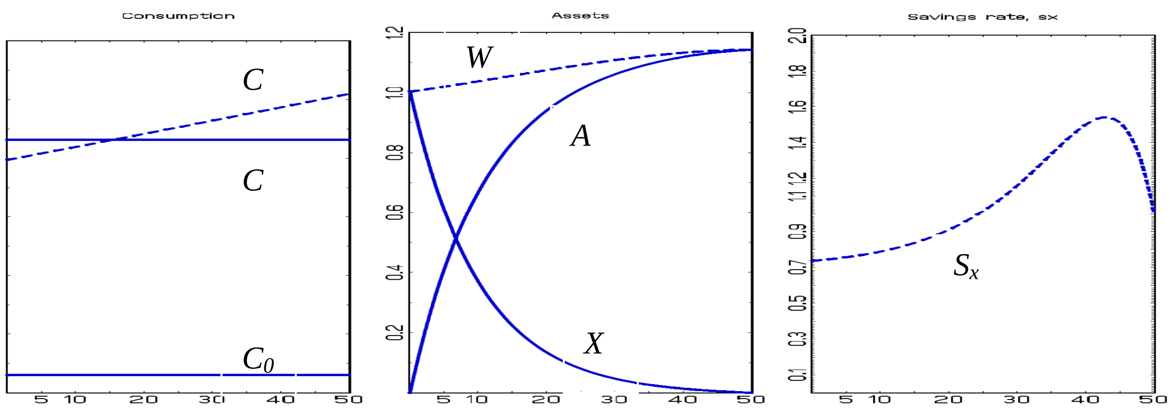

Allowing for uncertainty about the path of resource prices

So far we have largely abstracted from the path of resource prices, subsuming it into the concept of the path of resource revenues. A convenient starting assumption as to the path of prices is the Hotelling Rule, whereby rational speculation ensures that the price rises at the world rate of interest. However, this depends upon some heroic assumptions; three in particular are questionable. First, the component of supply of natural resources generated by actors whose decisions are based upon a rational expectation of their future price might not be large enough to determine the world price. Secondly, the costs of varying the rate of extraction might be too large to ignore, so that the extraction decision is not entirely determined entirely by the path of the expected future price. Thirdly, over the long term supply and demand might change due to innovations about which there is so little knowledge that no rational forecast can be made: actors face uncertainty as well as risk. Empirically, even over very long periods the Hotelling Rule shows little evidence of being a useful approximation to reality, let alone of being a sound basis for a government savings rule. Hence, it is worth considering the implications of relaxing the assumption that the rate of return on natural resources left in the ground must be the same as that on other assets. The government therefore faces uncertainty. If it errs on the side of caution, by assuming a rate of price appreciation which is lower than that which actually prevails, it will make systematic errors, on average being surprised that actual prices and revenues are higher than anticipated. In the limiting case, the government completely discounts all future resource revenues as being too uncertain to value. This is equivalent to persistently but erroneously assuming that future revenues will be zero. Each year the government therefore receives an unanticipated windfall which it treats as something that will not be repeated. In this case, the appropriate permanent income response to such a windfall is to save it all. This is the risk-based case for the 'Bird-in-the-Hand Rule'. The Rule imposes substantial misallocation of consumption and is thereby utility-reducing.

More generally, uncertainty can be represented as a discount upon future revenues. The present value of future expected revenues is $X_t \equiv \int_t^{\infty} x_\tau e^{-r_X(\tau-t)} d\tau$. Suppose that the rate at which expected future flows are discounted, r_X, is high, exceeding the rates of time preference and return on assets, $r_X > \rho = r_A (\equiv r)$. At each date the government follows the PIH, setting consumption equal to its permanent value given the value of wealth, $c = rW$. However, W is no longer constant. Using the differential equation for W given above together with $c = rW$, the evolution of wealth associated with the windfall is

$$\dot{W} = \dot{A} + \dot{X} = [r_A A + x - c] + r_X X - x = X[r_X - r_A].$$

The government is surprised in each period as outcome x_t occurs and W is larger than expected; $\dot{X}$ is greater than it would have been had the lower discount rate r_A been used. [2] If W is rising through time, then so too is consumption. Consumption starts low as future expected revenues are heavily discounted, and as a consequence savings rates and asset accumulation are high. W rises fast, and consumption comes to overtake its PIH level.

This is illustrated in figures 4a and 4b which look at the paths of consumption, asset accumulation, and the consequent saving rate under the pure PIH and when $r_X > r$. Figure 4a takes the boundary case in which resource revenues are declining so rapidly that the optimal savings rate is constant until near exhaustion when it finally rises towards 100 percent. We take the boundary case because we have established that in the more typical case of a rising optimal savings rate, uncertainty dampens the rise. By taking the boundary case we are able to check whether sufficiently severe uncertainty might lead to a declining savings rate.

The left hand panel in each figure gives the initial level of consumption (lower line), the PIH level (upper solid line) and, in Figure 4b the consumption trajectory when $r_X > r$, starting below and rising to overtake the PIH, as discussed. The second panel gives the path of X (solid declining line, assuming exponential decline in revenue flow), A (increasing dashed line) and total wealth, $W = A + X$. Total wealth is flat for the PIH, as increases in A exactly offset reductions in X. With uncertainty (4b) X_0 is small, not because realized x values are any different, but because of high discounting (the example has $\rho = r = 4\%$, $r_X = 20\%$). A and W both rise rapidly, supporting the rapid growth of consumption.

The third panel is the savings rate, s_X. For the PIH, this is as in the dashed line of figure 2, and, in this example is around 70%. When $r_X > r$ saving is higher at all dates, running at around 90% in the example. However, even the conjunction of a high discount rate for uncertainty with exponentially declining revenue does not cause the savings path to be downward-sloping. As previously, as exhaustion approaches, the path of the savings rate rises towards 100 percent. In the limiting case when $r_X \rightarrow \infty$, the savings rate starts are 100 percent and remains there.

[2] It is straightforward to add noise to x, making the paths of X and W stochastic.

Figure 4: Uncertainty: consumption, assets and savings

Figure 3a: PIH, $r_x = r$

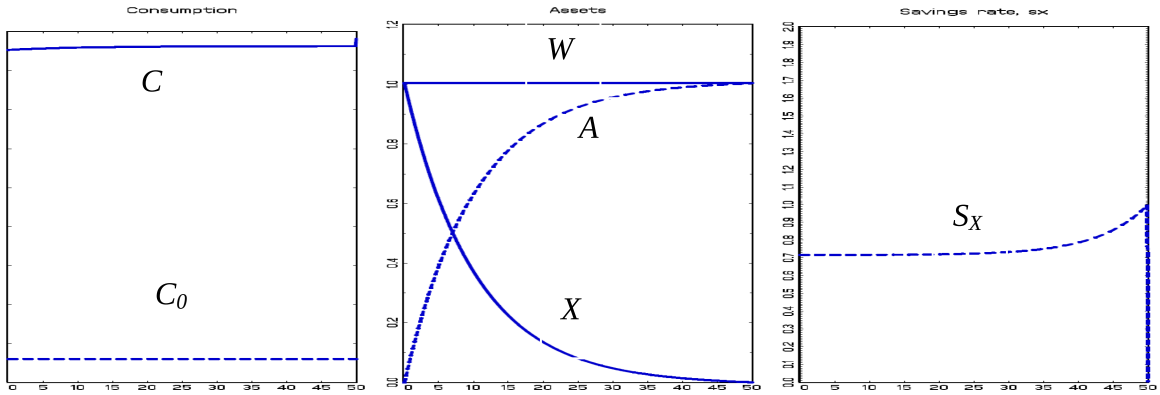

Figure 3b: $r_x > r$

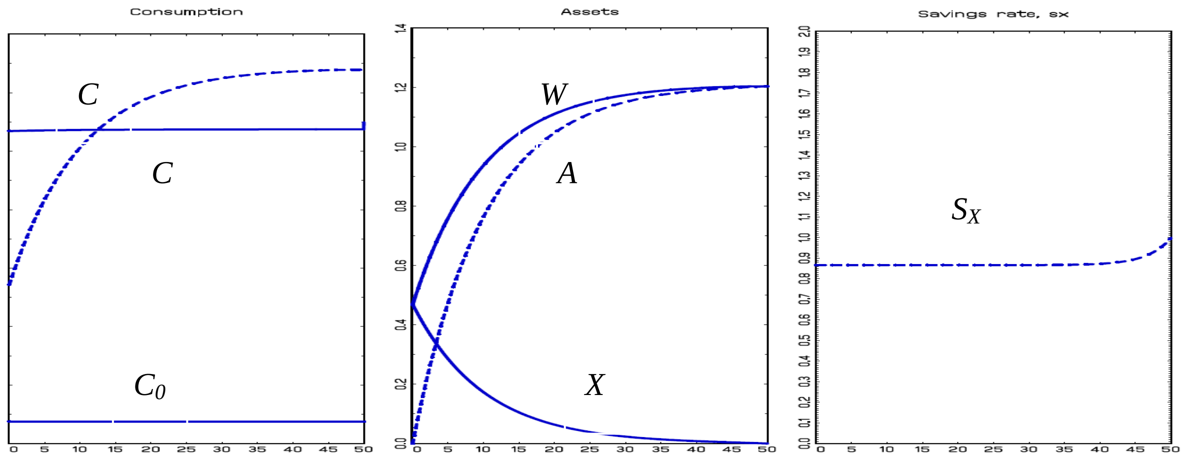

So Where Have We Got To?

We have now covered how savings should respond to the fact that resource extraction depletes the physical stock of the resource and so is unsustainable. We have abstracted from two potentially important issues: the path of physical extraction has been treated as exogenous, and volatility, as distinct from uncertainty of the path of revenues, has been ignored. The former we justify as being a reasonable approximation to the situation facing many developing countries. The latter, we will relax in the next section. However, it is useful at this stage to summarize the implications of our analysis and to contrast it with conventional policy prescriptions.

Our central result is that while the optimal saving rate just prior to exhaustion is close to 100 percent, in the first years of extraction it is likely to be much lower – in our examples typically around 30 percent. The optimal path of the savings rate is thus rising. This result is robust to all but the most extreme assumptions about extraction paths. Indeed, as the severe assumptions of the base case are relaxed, allowing for greater realism, the additional effects reinforce the result. Yet this central feature of the savings path appears to have been entirely missing both from the analytic and the policy literature.

Probably the most widely touted savings rule for resource-rich countries is for all resource revenues to be saved. The extra consumption afforded by natural resources is set equal to the revenues from the accumulating financial assets. This savings rule has two distinct intellectual underpinnings, one being management of uncertainty and the other the 'Hartwick Rule'.

The management of risk we have already discussed in the previous section. Uncertainty indeed increases the optimal savings rate, but the assumption needed to drive it to 100 percent is evidently extreme: all future resource revenues have to be discounted away to zero. Further, the risk that future revenues might fall short of plans does not exist in a vacuum: the typical poor society faces many other risks that stem from its low level of consumption. Indeed, Revallion (2012) shows that a high incidence of poverty in a society is itself grow-retarding. Hence, in this larger context, choosing to keep consumption low so as to reduce the risk that revenues will fall short may well be risk-increasing rather than risk-reducing. Imposing the infinite discount rate necessary to generate a continuous 100 percent savings rate would therefore be entirely unwarranted.

The Hartwick Rule, (Hartwick 1977), states that within the standard Utilitarian calculus used by economists to balance the interests of the future against the present, the extraction of depleting resources should be fully offset by the accumulation of capital. That is, a savings rate of 100 percent out of resource revenues is socially optimal. Because the Hartwick Rule was derived from a standard permanent income economic analysis and prestigiously published in the *American Economic Review*, it has long underpinned professional policy thinking on what constitutes ideal practice in resource depletion. Hartwick himself was Canadian, based at Queens, Ontario, and the original motivation for his analysis was how Canada should respond to its own resource depletion. Hence, it is unsurprising that it has been seen as providing intellectual underpinnings for the ideal savings rates out of natural resource revenues for resource-rich countries. However, since the original article was published over three decades ago, its continuing influence is now somewhat obscured beneath layers of more recent literature.

In fact, the Hartwick Rule as it has commonly been interpreted is highly misleading. The Rule derives directly from the Hartwick Model. Two striking features of this model immediately indicate why it is inapplicable to a resource-exporting country. First, it is a closed economy model: evidently, at best the model generates implications for the world in aggregate rather than for a resource-rich country exporting to resource-scarce countries. Second, although the model has three factors, capital, labour, and an exhaustible natural resource, it has only a single agent who owns all three factors and whose welfare is maximized. Hence, it is not a guide to maximizing behaviour for an agent who owns only the exhaustible natural resource.

Let us reconsider the Hartwick Rule, retaining the exact assumptions of the Hartwick Model with the necessary modification that one agent owns all the exhaustible resource, and another agent owns all the capital and labour, with the former agent selling the resource to the latter. For concreteness, we will refer to these two agents respectively as Africa and China.

Within the Model the natural resource is depleted in such a way that its value rises at the rate of return on capital, thereby replicating the Hotelling Rule discussed above. As this occurs, producers economize on its use so that the rate of extraction diminishes at an asymptotic rate. Indeed, given the Cobb-Douglas form of the technology in the Hartwick model, the rate of extraction diminishes at the same rate as the appreciation in the value of the exhaustible resource. Since depletion is precisely offset by appreciation, the value of the resource remaining in the ground is constant. In other words, the Hartwick Model assumes a special case of our asymptotic depletion case. As a result, Africa's sustainable rate of consumption out of the income from resource extraction is 100 per cent. Rather than needing to use all the revenues for accumulation as supposedly implied by the Hartwick Rule, according to the Hartwick Model Africa does not need to save at all.

The only context in which the Hartwick Rule is correct is for the world in aggregate. Whereas Africa continually gains from the appreciation of its natural resource, China continually loses from the decline in the marginal product of capital consequent upon its combination in production with diminishing quantities of the exhaustible resource. This inflicts a repeated capital loss on China's holding of the capital stock which, for sustainability, it needs to offset by accumulating more capital. Indeed, within the Hartwick Model, China needs to accumulate an amount of capital precisely equal to Africa's depletion of its natural resource. Thus, whereas Africa does not need to save, China, entirely lacking in natural resources, must adopt a high savings rate simply to sustain its level of consumption. Thus, precisely contrary to the common misinterpretation of the Hartwick Rule, the Hartwick Model implies that it is the resource-scarce economies that must save to offset resource depletion, not the resource-rich. The significance of this result is not that it provides new insight into the appropriate savings behaviour of a resource-rich country. The special case in which physical depletion is precisely offset by appreciation is so restricted as to be of no interest. Rather, its significance is that it demonstrates that the Hartwick Rule is irrelevant for the question of what might be the appropriate savings behaviour of a resource-rich country.

To summarize, neither uncertainty about future revenues, nor the Hartwick Rule, provide any justification for a savings path along which 100 percent of resource revenues are continually saved. The optimal savings path depends upon the specific details of the extraction path, but it will start with a savings rate below 100 percent, and most probably considerably so. The path will gradually rise, only approaching 100 percent in the final years prior to exhaustion.

However, depletion is only one influence upon the savings rate out of resource revenues. We next turn to volatility.

2. Savings in Response to Revenue Volatility

The above analysis solves for both optimal savings and for optimal consumption. Budgets, however, work with two different concepts, expenditure and revenue. Revenues equal the sum of consumption and asset accumulation (savings), but expenditures are the sum of consumption and domestic investment, a component of asset accumulation which we will discuss in Section 4. Because it is costly to deviate from planned optimal expenditure, changes in assets should accommodate deviations between it and actual revenue. Above we considered deviations arising from the expected revenue path. Now we consider deviations arising from periodic departures of actual revenues from expected revenues.

That path of future revenues is not known with certainty raises four distinct issues. One of these, uncertainty about the average rate of change of revenues, we have already discussed above: such uncertainty can be incorporated by an increase in the discount rate and this in turn raises the optimal savings rate. The expected path of revenues will need to be revised periodically in the light of new geological and market information and this will adjust the optimal path of savings and consumption. At the other extreme is short term uncertainty about prices: this is of particular importance because of annual budgeting. An annual budget incorporates, explicitly or implicitly, an assumption about the average price of the resource over the coming year and this assumption will inevitably prove incorrect. The third aspect of uncertainty is the volatility of the annual average of prices: even if, at the start of each year the average for the year were to be correctly forecast, there would be a need to react to the changes in the average between years. If annual expenditures are not to be volatile, recourse must be made to savings or borrowing. The final aspect of uncertainty is that periodically there will be unanticipated new discoveries. This we will defer to the next section where we consider transition more generally. Hence, this section will focus on the short term uncertainty facing annual budgeting, and the medium term uncertainty associated with fluctuations in the average revenue between years.

Short-Term Uncertainty and Annual Budgeting

Over a horizon of twelve months the path of physical extraction is largely known: discoveries made during the year are highly unlikely to be brought into production within this period. Hence, the main uncertainty concerns prices. However, for all significant commodities it is now possible to hedge prices over this horizon. The whole point of annual budgeting is to enhance the coherence of spending, and so there is value in reducing uncertainty over intra-year revenues.

Among hedging strategies, the first choice is in the form of payment. Either the payment can be explicit: a known expenditure to purchase a floor price, or the floor price purchased in exchange for a ceiling price. As between these, the former is likely to be preferable. Whereas it is important to avoid receiving a price below the floor price, there is no equivalent need to avoid particularly high prices since above some ceiling all revenues should be saved and marginal additions to such savings incur no cost. Hence, there is no point in paying an implicit risk premium to eliminate this range of uncertainty. Further, while the use of a ceiling may appear to have political advantages, disguising what would otherwise be an explicit budgeted payment, the circumstances in which it is triggered may be particularly damaging politically. Namely, an insurance against a very low price which has *a fortiori* turned out to be unnecessary has been paid for by sacrificing a high price which has materialized. A routine annual insurance premium for the purchase of a floor price securing the budget may be politically less exposed.

Having determined the form of payment, the remaining hedging choice is the precise floor price to be chosen. In the neighbourhood of the mean of market expectations, an additional dollar on the floor price will increase the cost of the hedge by around 50 cents. Hence, in this range half of the marginal revenue generated is eaten by the hedge: manifestly this is far too large a proportion to be warranted politically. An implication is that a floor price hedge should pitch the floor price conservatively, below the mean of market expectations. In fact, the floor price does not normally need to be particularly close to the mean of market expectations. The floor price is not itself a forecast but rather a way of protecting expenditure. Planned expenditure will on average be below expected revenue partly because not all planned savings will be invested domestically, and partly because expected revenues will be conservatively estimated due to a risk discount.

Hence, where the market expectation for the coming year is equal to the long run expected price, the hedged floor price needed to protect expenditure will be below this level.

Medium-Term Uncertainty and Inter-Year Smoothing of Expenditure

The revenues prevailing in any one year may be above or below the long-term risk-discounted, expected path. As long as actual (post-hedged) revenues are above planned optimal expenditure for that year then it is of no consequence for spending if they are below their expected level. All the difference between actual and expected revenues can be borne by a deviation of actual financial savings from planned financial savings. If, however, actual revenues are below planned optimal expenditure then either actual spending falls short of optimal, or the shortfall is financed.

In principle, finance can be through either borrowing or drawing down savings. However, in practice the two are often not alternatives. The ability to borrow depends upon a record of prudent savings and the prior accumulation of liquid assets. For example, during the global economic crisis of 2008/9 the government of Botswana was able to borrow $1bn in order to protect public spending but this was only possible because it had accumulated a much larger stock of wealth which it preferred not to draw down at such a time. For countries without such a record, the only reliable source of finance is the prior accumulation of liquid assets.

However, for developing countries, liquid savings have an opportunity cost in terms of more productive assets foregone. Hence, not all possible scenarios of needs for liquid savings should be accommodated: there will be times at which actual expenditure will indeed need to fall below planned optimal expenditure. As liquid assets are drawn down, actual expenditures should be pre-emptively reduced to avoid the risk that finance will be exhausted, forcing a large, abrupt reduction in expenditure. In effect, this override is a second line of defence against an overly optimistic assessment of the path of future revenues, protecting accumulated assets intended to offset depletion from being used to finance an unsustainable level of consumption.

If revenues exceed planned expenditure then the surplus should evidently be saved. However, there are two distinct functions for such savings: the accumulation of liquid savings to buffer expenditure, and the accumulation of longer-term financial assets as part of the strategy of offsetting the depletion of natural assets. Although in the long run a large majority of the assets that offset depletion should usually be domestic, reflecting the initial lack of domestic capital, the decision as to the composition between foreign and domestic investments should be taken year-by-year and reflect the limits on current capacities to invest well within the economy. Hence, in the initial years of resource revenues there is likely to be a substantial investment flow into foreign assets. There is therefore a need for some decision rule as to how much of the excess of revenues over planned expenditure should be used for future smoothing, and how much should be used for long term portfolio investment. As we discuss in Section 5, political considerations make it desirable to have distinct funds for these two purposes. One approach would be notionally to divide the surplus into that above and below expected revenue, with the former flowing into the smoothing fund and the latter into the depletion offset fund. An alternative is to decide the allocation according to whether the smoothing fund is at or below its target level. Until the fund reached this level all savings into foreign financial assets would be allocated to it, and beyond that all would be allocated to offsetting depletion. This second approach has the advantage that it avoids the anomalous situation of the government being required to add to its stock of foreign financial assets at a time when it would be justified in running them down. Such anomalies would not only be symptoms of misallocation; they might jeopardise an entire rule-based system of managing resource revenues.

Summary and Implications

To summarize, planned optimal expenditure is derived from the path of expected revenues, as discussed in Section 2. Having determined planned optimal expenditure for the coming year, these plans should be implemented even in the face of revenue shortfalls, subject to an override reflecting concerns over liquid savings. Because optimal expenditure is to be protected from revenue fluctuations, it is important that there is a responsible process of regular updating of expected revenues, so that planned optimal expenditure is based on realism tempered by risk. The override requires planned annual spending to be reduced below its optimal planned level if the country encounters a run of unexpectedly low revenues which drains liquid savings to a dangerous level, the danger being an enforced and abrupt reduction in spending.

The key operational concepts are the optimal level of expenditure, the average level of liquid savings to be held for maintaining expenditure at this level, and the rules for overriding the drawdown in savings. In turn, setting these parameters should rest on an analysis of the likely volatility of revenues, which is commodity-specific, and the likely costs of volatility in expenditure, which in turn will reflect specific features of the system of public spending. However, the entire process of managing volatility is likely to be second-order relative to the task of saving to offset depletion.

Short-term price volatility is liable to upset annual spending plans even once they have been adjusted to reflect concerns about liquid savings. However, intra-year volatility in expenditure can be greatly reduced by hedging. We have suggested purchasing a floor price sufficient to cover planned optimal expenditure, this normally being below expected revenues.

3. Savings and Investment in Response to Transition

The two previous sections have essentially focused on equilibrium responses to a resource discovery. However, each involves a transition to the new equilibrium: the news of the discovery makes the prior equilibrium sub-optimal. In Section 2 we showed that the discovery will involve a quantum increase in equilibrium expenditure. Yet expenditure cannot efficiently be increased in a quantum manner: both decisions and implementation require a gradual adjustment to the new level. Similarly, in Section 3 we suggested that for a given level of revenue volatility and costs of expenditure volatility, a country would need to accumulate a particular level of liquid savings. Yet typically countries do not start with such savings: indeed, to have been holding large liquid assets prior to the news that they would be needed because of a resource discovery would have been sub-optimal. Hence, the country needs to determine an adjustment path along which these savings are accumulated.

The Transition in Expenditure

In considering the transition to a higher level of expenditure, it is necessary to distinguish both between consumption and investment, and between internationally tradable goods and those which are internationally non-tradable. For ease of reference, we will denote tradable consumer goods as 'manufactures', non-tradable consumer goods as 'services', tradable investment goods as 'equipment', and non-tradable investment goods as 'structures'.

The adjustment between the initial and the eventual equilibrium involves a larger percentage increase in investment than in consumption. This follows because the overall savings rate for the society will be rising, and being capital-scarce it is appropriate that eventually most of the extra savings is used to finance domestic investment. Increases in demand for tradable goods, whether for investment or consumption - equipment or manufactures -, can be supplied by imports. It is reasonable to assume that resource-rich developing countries are price takers on these global markets. The potential for supply constraints thus only arises in respect of non-tradable investment goods – structures, and non-tradable consumer goods - services. Of these, supply constraints in structures are more important because as noted, the increase in demand should be proportionately greater. Structures are supplied by the construction sector, and so the elasticity of supply of this sector is the key potential bottleneck during transition. A range of disparate policies can potentially increase this supply elasticity: for example, the legal framework for land acquisition, and training programs for construction skills. Such microeconomic interventions have an important bearing upon macroeconomic choices because they determine the feasible pace at which the demand for structures can be increased without dissipating additional investment expenditures in increased unit costs.

In turn, the pace at which the supply of structures can be increased is one constraining factor on the rate at which investment can be scaled up: while there is some scope for increasing investment in equipment ahead of structures, the two forms of capital are evidently complements and so cannot get too far out of line. However, it is not the only constraint: the capacity to design, select and implement investment projects can only be increased gradually. Addressing these constraints in the construction sector and in project management is a necessary prelude to the actual scale-up of investment; it can be thought of as 'investing-in-investing'. This phase does not require major expenditures; rather, it requires policy attention. As such it is well-suited to the period between the news of a resource discovery and the actual extraction of resources and generation of revenues. A typical lag between discovery and revenues is around five years, and this is long enough for substantial progress in investing-in-investing.

Once the capacity to invest has been enhanced, then the next step in the transition to a higher level of consumption is to scale up investment. Because extra consumption of services, like extra investment in structures, must be met by increased domestic supply, the priority is for investment to be targeted on these sectors. Hence, precisely counter to prevailing policy concerns about diversification of exports, investment should initially be skewed towards the non-tradable sectors.

Only once investment in the services sector has increased can consumption be increased without forcing up the relative price of services – that is, without incurring Dutch disease. However, even then, consumption cannot leap instantly to its new level. As with investment, additional consumption expenditures must be designed, selected and implemented and this requires managerial capacity. The period of transition is, however, likely to be shorter than for investment, both because the proportionate increase in consumption is smaller, and because consumption typically requires less implementation capacity than investment.

In summary, the transition sequence towards the permanently higher level of consumption warranted by a resource discovery has three phases. The first phase is 'investing-in-investing' – augmenting the capacity to invest well. The second phase is to scale up expenditure on investment. The third phase is to scale up consumption. This sequence has implications for the optimal rate of saving. Essentially, since the increase in consumption should be delayed, there is a case for a temporary increase in the savings rate. Quite how long this is depends upon how long it takes to invest in investing, and to scale up investment, and when these processes can get underway.

If, as suggested above, the investing-in-investing phase can be accomplished between the news of a discovery and the onset of resource revenues, then the duration of temporarily high savings is correspondingly shortened.

Once the capacity to invest has been sufficiently improved, increased spending on investment should precede increased spending on consumption, but since the delay here is only due to the lag between investment and the output it generates, the lag need not be very long: investing-in-investing is likely to take much longer than the time needed for investment to generate output. However, since investment must proceed consumption in order to moderate Dutch disease, this creates a potential counter to the principle of a rising savings rate (and hence declining consumption rate) enunciated in Section 2. Evidently, given the lag between investment and extra output of services, the society cannot begin resource extraction with a high consumption rate out of revenues.

Depending upon the lag between the news of discovery and the onset of substantial revenues, these may be a case for bringing forward investment in advance of revenues. For example, if the news of the discovery provides six years advance warning of the onset of revenues, then the first four years might be devoted to building the capacity to invest, while the last two are used to start scaling up investment financed by borrowing. A policy that permitted borrowing on the collateral of future revenues to finance investment but not consumption might be politically viable, reconciling urgency with prudence. In this case, the pre-revenue phase could accomplish the transition in investment. Consumption expenditure should still not leap on the onset of revenues, but the problem may be less acute than it seems because revenues typically build up gradually, so that a rule of no borrowing for consumption may be sufficient to enable the revenue phase to be governed by the principles enunciated in Section 2, consumption expenditure in the early years being limited simply by low revenues rather than by a temporarily high savings rate.

The Transition in Liquid Assets

The country starts with insufficient liquid assets to smooth expenditures. It therefore needs to accumulate them up to a target level. Above we have suggested an equilibrium rule whereby at times when the liquid fund is below its target level all of the excess of actual revenues over planned expenditures is allocated to replenish it. This rule can be extended to cover the initial process of building the fund up, subject to a proviso. The difference between replenishment in equilibrium and initial accumulation is that the former is precisely matched by periods in which all surplus money is allocated to offset depletion whereas the latter is not. Hence, the replenishment rule does not imply any net reduction in savings to secure permanent income, whereas applying the same practice to the phase of initial accumulation might potentially transfer revenues from saving to consumption. The most straightforward solution to this potential problem is to assign the value of the smoothing fund at the point of resource exhaustion to the depletion fund. In this case, the expected end-value of the smoothing fund is part of the depletion fund and so the revenues used for its initial accumulation are equivalent to revenues placed directly in the depletion fund. Income from the smoothing fund once it is at its target level can therefore be treated as permanent income and so, like that from the depletion fund, used to finance consumption.

4. From Analytic Principles to Policy Rules

Budget rules for resource-rich countries

Budgets in resource-rich countries are essentially decision processes for allocating revenues to a variety of expenditures. Over recent decades two principles of good budgeting have been widely accepted; one macro the other micro. The macro principle is that aggregate expenditures should be kept broadly in line with aggregate revenues (which for present purposes we will treat as exogenous). Commonly, this principle has been encapsulated in the *balanced budget rule* which sets ceilings both for the fiscal deficit and for the ratio of debt to GDP. The micro principle is that the marginal benefit of expenditures should be equated across categories. Since it is reasonably assumed that priorities will change over time, this principle has been encapsulated in the *integrated budget rule* which discourages pre-commitments of revenues, and earmarking of particular revenue streams, to particular items of expenditure.

Resource-rich developing countries face distinctive fiscal problems, the solutions to which involve distinctive principles. They therefore need distinctive rules which encapsulate these principles. At the core of the distinctive problem is that, unlike other countries, the key revenue source for government spending is *unsustainable*. Manifestly, this is not a marginal problem: unless it is properly faced the consequences for welfare are liable to be disastrous. Hence, we should not expect it to be adequately dealt with by minor tweaks to the budget process. It is particularly problematic because inter-temporal resource allocation is an issue that standard budget procedures barely address: budgets are essentially devices for annual commitment. At the most, governments announce Medium Term Fiscal Frameworks, but these are essentially informal statements of intentions over a three year horizon; not only do these statements of intent have only limited credibility, there is no intention that they will bind a successor government beyond an election.

The distinctive principle for a resource-rich country is that a certain proportion of revenues should be saved, whether in financial assets or domestic investment. Conventional budgeting processes are inadequate to deal with this problem in two important respects. First, as noted above, they lack any mechanism for inter-temporal commitment, most especially over the horizon of around a generation which is likely to be required for resource depletion. Commitment technologies are valuable to governments to reduce the risk of temporary lapses resulting from random short term political pressures: that is indeed why they adopt the balanced budget rule and the integrated budget rule. But in the case of savings out of resource revenues a commitment technology is even more important. It is not just that without it there is a risk of a random lapse; rather, without it the incentive to save is reduced even for a good government. Without a commitment mechanism, the savings of one government may merely transfer spending power to a bad successor. Indeed, the rationale for augmenting permanent income depends upon the current government believing that *all* future governments will behave prudently. In the absence of a rule, a good government may reasonably decide that it is better to spend all the revenue now on items that it regards as desirable, rather than risk its savings being spent by a successor government on items that the current government regards as less valuable. Worse, without a commitment technology, as wealth accumulates the incentive to be a rogue government that favours only expenditure on consumption actually increases. Hence, a long-term savings rule is not a mere nice-to-have addition to the standard budget rules, it is paramount.

Second, because the balanced budget rule is defined in terms of expenditure relative to revenues, it misses the key required distinction between expenditure on consumption and the acquisition of assets. Domestic investment, which is aggregated under the balanced budget rule with consumption as expenditure, is the activity which for a resource-rich country it is most important to distinguish. Hence, the recent practice of modifying the balanced budget rule so as to exclude resource revenues – through concepts such as 'the non-oil fiscal balance' – has no analytic basis. Indeed, the government of a resource-rich developing country which actually constrained total expenditure to be equal to non-resource revenue would be massively misallocating its resource revenues, both under-consuming and under-investing. Conventional budget rules cannot be restored to relevance simply by setting resource revenues to one side. Rather, the principles underlying the optimal management of resource revenues must be woven into the foundation of a distinctive set of budget rules.

From Analytic Principles to Practical Budget Rules

It is now time to pull together the analytic principles enunciated in Sections 2, 3 and 4.

Rules for Offsetting Depletion

Section 2 set out the principles of the asset accumulation appropriate to offset depletion for an initially poor country. There were several distinct analytic effects. The *depletion effect,* which encapsulates the changing balance of natural and accumulated assets, unambiguously implies that the savings rate should rise over time, reaching 100 percent at the point of exhaustion. As exhaustion approaches this effect always comes to predominate. The *extraction path effect*, which allows for a non-constant path of resource revenues, modifies but does not fundamentally change these results. The most likely profile of resource revenues is hump-shaped, and as illustrated in the right hand panel of Figure 2, even with such a complex pattern of rising and then declining revenues, the savings rate nevertheless is monotonically increasing over time. The boundary conditions for when the savings rate would not be monotonic are extreme: resource revenues would need to be declining more rapidly than exponentially. The *divergent discount rates effect* incorporated the impatience likely in a poor society surrounded by a world of plenty: allowing the domestic rate of discount to be above the world discount rate (which was the rate pertinent for the Hotelling Rule). This effect also produced a rising savings rate. The discovery of natural resources is likely to raise the return on investment in the economy, pushing the capital market into temporary disequilibrium as capital is accumulated. We showed that this *capital markets disequilibrium effect* also led to a rising savings rate. Finally, we introduced uncertainty as to the extraction path: in response future revenues should be discounted. This was the only effect which tended to offset the tendency of the savings rate to rise. However, we showed that it could never do more than flatten the savings schedule: it could not generate a phase in which the optimal savings rate is declining over time. Even the case of a completely flat schedule only arose when the savings rate was 100 percent throughout. This in turn required the extreme response to uncertainty in which all future resource revenues were discounted to zero, whereas all the risks associated with continued poverty were dismissed. Hence, the sum of these various effects is unambiguous: the optimal savings path rises during depletion, ending at 100 percent. As we illustrated in Section 2, the issue is not merely a matter of fine-tuning. In the first years of depletion the optimal savings rate might be only around 30 percent. Any constant savings rate rule will lead to substantial inter-temporal misallocation of expenditure.

How best might the principle of a rising savings rate be incorporated into rules? Formulating a rule involves a trade-off between accuracy and simplicity. The less accurate is the rule the more sub-optimal the allocation it will generate and the more subject is it to challenge and change. However, it is more difficult to build a critical mass of citizen support for a complex rule than for a simple rule. While ignoring the first derivative of the savings schedule – the fact that it is upward sloping – would impose major costs of misallocation, ignoring the second derivate – the fact that it is unlikely to be linear – is probably warranted. Hence, a sensible compromise between accuracy and simplicity may be to have a rule in which the savings rate starts at some modest rate, rising annually through the lifetime of extraction. For example, the revenues from an oil discovery with an economic life of 25 years might start with a savings rate of 25 percent, and rise annually by three percentage points so as to end at 100 percent. Would such a savings rule be practical politics? Arguably, it might be more practical than a rule for a constant savings rate that cumulated to the same asset value. Evidently, it is easier for politicians to commit to the formula 'God make me good, but not yet'. This is, indeed, the explanation for the infamous 'weeping willow' pattern of medium-term budget projections: the government acknowledges that this year spending will rise, but reassures markets that this will be reconciled by future fiscal tightening. Yet in respect of savings from natural resource revenues, such a behaviour profile is actually optimal: the politically easy is the appropriate decision.

What are the practical alternatives to *the rule of a rising savings rate*? How might greater complexity be incorporated? One approach would be to acknowledge irreducible complexity and park the decision on the savings rate with an institution properly equipped to handle complexity. In effect, a critical mass of citizens would need to understand the limits to their own understanding. There are, of course, many precedents for such acceptable specialist authorities, for example Supreme Courts which interpret complex constitutional laws. In respect of natural resource revenues Chile has come closest to this model of delegating complex decisions to specialists, although its institutions are as yet too new to have demonstrated robustness. A second approach would be to build the aggregate savings rate from discovery-specific savings rates: each discovery would have an optimal savings path determined by its estimated extraction path. This has the advantage of responding to fresh information, but risks being so complex that parliament cannot effectively scrutinize whether the savings rule is being followed. A third approach is to adopt the rule of a rising savings rate, but subject to periodic review (such as every five years) by a specialist and trusted agency able to understand complexity, such as a central bank. If, for example, major discoveries have fundamentally changed the extraction path, the savings rule might be revised. A further advantage of building in such a review process is that it may help to pre-empt opportunistic challenges, analogous to the way in which a schedule of elections protects a democratic government during a phase of unpopularity.

Rules for Managing Volatility

Section 3 set out the principles for smoothing expenditures in the face of revenue volatility. We proposed four such principles. One straightforward principle was the *hedging rule* which proposed that governments should lock into budget assumptions by hedging resource revenues for the forthcoming year. A more complex, but fundamental principle was that governments should make an assessment of optimal expenditure: that level above which revenues should be parked, and below which expenditures should be sustained by drawing on liquid assets. We emphasized that it was important to ground this estimate in realism. One approach is the Chilean panel of independent experts. Another is to adopt a mechanical rule such as a long-term moving average of commodity prices. The mechanical rule has the advantage that it cannot be as readily manipulated by political pressure and is more transparent. Further, it is manifestly not a forecast and so it is easier for citizens to understand that it is liable to be wrong and so needs an override rule as a second line of defence. That second line of defence is the third proposed rule, namely that if liquid assets become dangerously depleted due to a run of misfortune, expenditure needs to be reduced pre-emptively below its optimal level rather than risk an abrupt collapse in expenditure upon the exhaustion of liquid assets. A simple and effective formulation for the rule is that in no year may more than a certain proportion (such as a quarter) of the remaining liquid assets be withdrawn. In the event of persistent over-optimism in revenue projections this imposes a gradual adjustment to reality. The final rule specifies the target level of liquid assets. This is analogous to the conventional target for foreign exchange reserves, commonly specified as so many months of imports. In the case of resource revenues, the numeraire should evidently be the revenues themselves rather than imports: hence the rule would be that *liquid assets for purposes of smoothing revenue volatility should be built up to a certain multiple of resource revenues*. The actual multiple can only be determined by studying the expected volatility of revenues and the damage that expenditure volatility would inflict.

In equilibrium the government will thus be holding foreign financial assets for two different purposes, expenditure smoothing and offsetting depletion. There is a good case for holding these assets in separate funds, with distinct rules. The key difference is that the depletion offset fund needs to be protected from being plundered to finance consumption. As discussed above, the distinctive challenge facing resource-rich societies is the need for commitment to long term asset accumulation. Hence, this needs to be embodied in the rules of the depletion offset fund. The rules cannot, however, be as simple as saying that financial assets cannot be liquidated, since it will be optimal gradually to shift the composition of the fund from foreign financial assets to domestic investment. Hence, the rule should be that *assets cannot be liquidated to finance consumption*. To distinguish this accumulation of both foreign and domestic assets from a conventional Sovereign Wealth Fund, we refer to it as a *Sovereign Development Fund*. In contrast, this is the part of the legitimate purpose of the smoothing fund. It is not the only purpose, since the fund is meant to smooth all expenditure, both on consumption and investment. Hence, for this fund, the rule should be that the assets cannot be liquidated to finance expenditures in excess of the planned expenditure. Further, it might be useful to place the implementation of the hedging rule within the smoothing fund: analytically, it is a means of achieving expenditure smoothing, and politically it enables the Finance Minister to be distanced from the decision to spend money on any particular hedge which may or may not turn out to have been vindicated by events. The purpose of the fund is thus to make expenditure resilient to revenue shocks, whether by hedging or by the accumulation of liquidity, and so we term it a *Sovereign Liquidity Fund*.

Rules for the Transition

Section 4 discussed two principles of transition, one concerning consumption and the other the stock of liquid assets.

The principles for the transition to the new equilibrium level of consumption proposed a sequence: investing-in-investing, then gradually increasing domestic investment to its newly optimal level, and finally increasing consumption to its newly optimal level.

Underpinning investing-in-investing is adherence to good practices in the design, selection and implementation of investment projects, and a range of policies to reduce the unit costs of capital, and to improve the policy environment for private investment. The details of these investment-focused policies would take us beyond the purpose of this paper. However, they suggest that the cluster of rules concerning these investment processes might come under the remit of a specialized agency such as an Investment Authority.

Our concern here is therefore with the implications of the above sequence for the deferral of consumption until after investment has been increased, and the deferral of investment until after a phase of investing-in-investing. How might simple rules encourage such deferrals? We have suggested that because the news of discovery precedes revenues by several years, this may provide a natural opportunity for sequencing: investing-in-investing does not actually require significant new expenditure. Indeed, the modest financial requirements of building the capacity for investment could be provided through aid. The key issues are therefore to avoid bringing consumption forward ahead of revenues, while perhaps enabling investment to be modestly increased somewhat ahead of revenues. We have suggested that this can be approximated by a rule which prevents borrowing on the collateral of future resource revenues except for spending on investment. A refinement might be severely to limit the maturity of such borrowing, thereby pre-committing some of the revenues in the early years of extraction to be used to repay money borrowed to finance prior investment. In 2011 the government of Ghana adopted a constitutional rule with some of these features, so the political economy may well be feasible. By reducing the revenues available for expenditure in the first few years of extraction, this would help to limit the pace at which both investment and consumption were increased. Hence, the proposed transition rule is *to limit borrowing to the purpose of investment, and with a short maturity*.

The transition to the optimal stock of liquid assets proposed a rule that whenever the stock was below its target level (that is a certain multiple of resource revenues), all the surplus of revenue over planned expenditure should be devoted to replenishing it. Conversely, once the stock had reached its target level, all surplus revenue would be devoted to the accumulation of long term assets. Finally, at the point of resource exhaustion, the value of the smoothing fund would accrue to the depletion fund, thereby ensuring that savings used initially to accumulate liquid assets for smoothing were ultimately not a diversion from savings to offset depletion.

Comparing the Rules with Conventional Budget Rules

How do these rules align with the conventional budget rules: the balanced budget rule, and the integrated budget rule?

First, consider the balanced budget rule as applied to a resource-rich country. In its unmodified form it would preclude the accumulation of assets other than domestic investment, and so seriously distort the process of asset accumulation. In its modified variant of the 'non-oil balanced budget' it has the opposite distorting effect of squeezing out domestic investment. More trivially, the balanced budget rule collides with the need to smooth expenditure: self-evidently, resource-rich countries need rules for enabling, and indeed requiring, expenditure to deviate from revenues. Finally, as discussed above, the transition may be aided by a phase of borrowing followed by a phase of debt repayment. Hence, the balanced budget rule is irretrievably inappropriate for a resource-rich developing country: it is not up to the central task of inter-temporal resource allocation.

Now consider the integrated budget rule: the principle that all expenditures should be left uncommitted so as to be freely allocated each year. As will now be apparent, this is also fundamentally at odds with the need to pre-commit some revenues to asset acquisition. Without such pre-commitment there is little chance that the marginal equivalences between expenditures, which are normally the ultimate justification for an integrated budget, can be maintained.

Were the government to have full information about all future needs and revenues, the optimal budget process for a resource-rich developing country would not be a series of annual budgets, but rather a single inter-temporal budget over the horizon of resource depletion. Such a budget would incorporate the optimal path of asset accumulation, thereby achieving the marginal equivalence between the value of current and future consumption. Obviously, no government has full information and so such a comprehensive budgeting process is inappropriate. The solution is to leave open as many expenditure decisions as possible, locking in only to the minimum necessary to ensure the inter-temporal equivalence of consumption expenditures. This is what is achieved by the rule of the rising savings rate. Having pre-determined savings, the composition of savings as between domestic investment and financial assets can be left open to the annual budget, as can the allocation of consumption spending between items. If we conceptualize all present and future uses of revenue as a matrix, with the rows being the years and the columns the various uses, the annual budget pre-commits the current row, while the rule of the rising savings rate pre-commits the assets column. All other items are left for future decision. This structure mirrors the pattern of markets, with the markets for goods largely confined to the present period, while transactions concerning future periods are accommodated in an aggregated form through the capital market. Hence, the rule of the rising savings rate replaces the balanced budget rule and introduces a constraint into the integrated annual budget rule.

5. Conclusion

The management of new resource revenues is going to be the single most important economic challenge facing low-income countries. The standard rules of economic management were not devised for such countries and so fail to address problems that are distinctive but central. There are no suitable models to follow. The emerging market economies provide a highly successful model for development, but very few of them are resource-rich and so in core respects that model would be misleading. The handful of resource-rich high-income countries, such as the Gulf States, Norway and Australia, are structurally so different from low-income countries as to be seriously inappropriate as models. Hence, the governments of resource-rich low-income countries are on their own, facing issues that they are ill-equipped to deal with due to weaknesses in analytic and managerial capacity.

In this paper we have focused on one distinctive and central policy issue, how much of the revenues from natural resources should be saved? We have combined the three core features which make resource revenues distinctive: they are depleting, they are volatile, and the news of their discovery creates a quantum leap in wealth and so plunges the economy into disequilibrium. The principles concerning the optimal responses to these features have not been well explored by the established literature. Although in this paper we have elucidated these principles, we have not taken the analysis through to country-specific applications: for this simulation modelling would be necessary. Finally, we have suggested how the principles can be approximated into practical policy rules. An implication of the paper is that the fiscal rules that have become conventional for countries that are not resource-rich are seriously inappropriate. The issue cannot be addressed by minor tweaking of conventional fiscal rules. The challenges that resource-rich, low-income countries face are sufficiently distinctive and critical for them that purpose-designed fiscal rules are necessary.

FERDI

Créée en 2003, la **Fondation pour les études et recherches sur le développement international** vise à favoriser la compréhension du développement économique international et des politiques qui l'influencent.

Contact
www.ferdi.fr
contact@ferdi.fr
+33 (0)4 73 17 75 30