

Resource Revenue Management: Three Policy Clocks

Paul Collier

Blavatnik School of Government,

Oxford University

and

The International Growth Centre

October, 2014

Introduction

Economies in which the extraction of a non-renewable natural resource is a significant activity pose two distinctive challenges for economic policy. Revenues are likely to fluctuate because commodity prices have historically been volatile. Further, the revenue from extraction is generated by depleting a finite resource, and so there is potentially a case for offsetting depletion with the accumulation of other assets. Volatility and depletion work to radically different timescales: managing them evidently requires distinct ‘policy clocks’. Chile has led the world in its approach to managing volatility, but it has yet to think through the issues posed by depletion with equivalent rigor. Hence, my initial focus will be on whether Chile should be at all concerned about depletion, and if so, what an appropriate policy response might be.

Policy Clock 1: Offsetting Depletion

The conventional framework for thinking about the depletion of a finite natural asset is permanent income (PIH). While this framework is inadequate, it is a useful starting point. A proposition that follows immediately within this framework is that *the longer is the duration of the resource the lower should be the savings rate*. If the resource lasts for only one year, the optimal savings rate will approach 100 percent, whereas if it is expected to last for a century, it can be far lower. Hence, the issue is how long will Chile’s copper last.

Posed as a question of physical supply, the evident answer is ‘a very long time’. However, while the physical limits of availability are pertinent in some contexts, such as the discovery of a small oil field, for Chile a more pertinent issue is long term price risk. In predicting the long term price path of a natural resource there is a seductive temptation to invoke the Hotelling Rule. Shortly I will explain why I think this would be a mistake, but first, for those who believe Hotelling, I will set out its rather dramatic implications for savings. The Hartwick Rule (Hartwick, 1977) purports to answer this question. According to its analysis, in order to maintain living standards constraint, all revenues from a depleting resource should be devoted to asset accumulation. However, although this answer is usually taken to appertain to a resource-exporting country, in fact it applies only to the world as a whole, and absolutely not to a resource-rich country that sells its output to the rest of the world. Indeed, once the Hartwick model is decomposed into a resource

exporter (Chile) and a resource importer (China), its assumption of unit-elasticity of demand has the consequence that the gradual reduction in export quantities is precisely offset by the rise in relative price. As a result, for the resource exporter, revenue is permanently constant so that within the PIH framework it should save nothing. All the global saving effort required by the Hartwick Rule is done by the resource importing country, China, because it is paying a constant total amount for a diminishing physical flow.

I would not advise Chile to rely on the Hotelling Rule. Distant but major technological advances cannot be forecast. Almost certainly their prospect does not strongly influence current extraction decisions, yet this is the essential assumption for the validity of the Hotelling Rule. But distant future technological discoveries might generate structural breaks in the path of future prices. There are two such risks, one concerning demand and one supply. Much of the demand for copper is due to its properties of conductivity. It seems to me quite likely that by 2114 electricity will no longer be transmitted through copper wires. On the supply side, the world has virtually infinite amounts of copper: Chile's dominance of the global copper industry is entirely because on present extraction technology its endowment is considerably cheaper to mine than elsewhere. As extraction technologies change this advantage may persist, but then again, it may not: currently, Chile has an extraordinary 27 percent share of global copper production. The probabilities of these types of technological advance cannot reasonably be assessed: the typical horizon of likely technological advance seems to be around three decades (for example, revolutionary change in nuclear technology has been forecast as being four decades away for the past half century). These distant technological uncertainties are no different for resource extraction than for any other industry: any national industry might become unviable due to technological change. The reason why such change is uniquely important for the resource extraction sector is because of the consequences: far more is at stake. Resource extraction generates substantial rents whereas industrial sectors do not. Technological change in manufacturing and services – creative destruction – merely shifts the activities where capital can earn a normal return. Whereas countries that are fortunate to have major deposits of high-value natural assets enjoy substantial rents over-and-above the normal return on capital and it is these rents which are at risk from sector-destroying technology. Initially, Chile's post-socialist policy on copper was based on the assumption that there were no significant rents in the activity. The

large and prolonged increase in global prices of natural resources has required that assumption to be revised. Currently, and prospectively, the copper sector generates substantial rents, much of which accrue in Chile: some to the government, some to the national copper company, and some to workers in the sector, and local companies supplying it.

An analogy is the market in oil, which until recently seemed the best candidate for the Hotelling Rule. Four decades ago global warming was not part of serious policy discussion. Yet we now know that we have already discovered more oil than can be safely burnt. At some stage in the next half-century, rules are likely which will result in ‘stranded assets’ – potentially destroying all the rents on oil. The first indications of this risk are already apparent: major portfolio investors are beginning to divest from oil companies. As expectations adjust to this scenario, eventually oil prices will start to trend down rather than up consequent upon the ‘green paradox’ - a race to extract.

I have no expertise in forecasting the price of copper fifty years hence. You may conclude that the rents that Chile currently gets from its copper sector can be relied upon so far into the future that depletion can be ignored, or you may decide that faced with substantial long-term downside uncertainty it is better to be prudent. I will proceed on the assumption of prudence, since this seems the more appropriate stance for a central bank. In effect, were you to regard the rents from copper as secure only for the next three decades, what would be the implications?

The obvious high-level implication is that some of the resource rents should be used for the accumulation of long-term assets. This is quite distinct from the need to smooth short term fluctuations in rents, which Chile already manages. Beyond this general case for asset accumulation even the simple PIH model with which I started can offer more specific guidance.

Within the PIH model, the consumption attributable to resource extraction must be constant at its permanently sustainable rate. But what is seldom appreciated is that this implies that *the savings rate from resource revenues should rise over time*. The reason for this is that as the stock of assets resulting from the savings accumulates, the income generated by this stock, which is permanent and so can be devoted entirely to consumption, increases. Indeed, by the point at which the resource is exhausted, by definition this sustainable income stream must equal the

consumption target. Hence, with sustainable consumption almost fully financed by the income from accumulated assets, the savings rate from the revenues generated by the last resources extracted should be virtually 100 percent. In contrast, the savings rate from the first resources extracted should only be revenue minus target consumption. While this is analytically straightforward, it has not been widely noticed, and the practice of committing to a rising long term savings rate out of resource rents has not yet been incorporated into public policy in any resource-rich country.

While the assumptions of the above simple case were stringent, the rising savings rule turns out to be almost-general. Only if the path of extraction were to generate an extremely rapid decline in revenues (which is highly unlikely) would the optimal savings rate not rise. In the more normal case in which the extraction path is hump-shaped, a first phase generates rising resource revenues, followed by a second phase in which they are declining. During the first phase, the rule of the rising savings rate is reinforced: savings can increase because a smaller proportion of revenue is needed to maintain the PIH level of consumption. During the second phase this effect goes into reverse, but because the point of exhaustion is now closer, the annual effect of its approach is powerful. Hence, as noted above, only a very rapid decline in revenues (more than exponential) will overpower the need for a continued gradual increase in the savings rate.

While the PIH framework is a useful starting point, it is not adequate for an economy such as that of Chile which is still only about halfway to converging on the level of income found in fully developed economies (Van der Ploeg and Venables, 2011). Essentially, people are poorer now than they will be in the future. Somewhat offsetting this, during the convergence phase the rate of return on capital should be atypically high. The net effect of these opposing forces within a Utilitarian framework is to increase the optimal initial consumption level relative to that under PIH assumptions, but to gradually taper the consumption increment. This provides a further reason for the savings rate to rise during extraction.

The practical import of the above discussion is that a prudent policy would perhaps regard copper rents as secure for only the next three or four decades, while trusting that by the end of the century Chile will have fully converged on the OECD levels of income. By the time in mid-century that the rents might be seriously dwindling, the appropriate savings rate out of them

would therefore be high, likely above 50 percent. Currently, long term savings out of resource rents are essentially no different from any other source of income, and this makes no allowance for their distinctively uncertain sustainability. Hence, an initial extra savings rate might be set at a few percentage points over the average for all other income, while the rising savings rate rule would gradually raise this during the ensuing three or four decades. Note that this would apply to the rents on copper extraction, over-and-above the normal return on the capital deployed in the sector.

Policy Clock 2: Managing Asset Accumulation

Trivially, the assets held to offset depletion should differ from those used to smooth expenditure in the face of fluctuations in revenue. The latter must necessarily be foreign assets otherwise they cannot smooth domestic activity. Since they are being held in order to be liquidated when needed, they must be readily marketable. Further, since they will need to be liquidated in predictable circumstances – namely a fall in the copper price – they should be chosen so as to have a marketable value that is negatively correlated with the copper price.

In contrast, the assets accumulated to offset depletion are held for their long term return rather than for their ability to smooth domestic activity. The key issue for assets designed to offset depletion is the choice between investment in foreign financial assets and domestic real assets.

Analogous to the PIH framework, the simplest starting point is an economy fully integrated into global financial markets. If there is diminishing marginal productivity then incremental wealth should be held entirely in foreign assets. Increasing the capital stock would reduce the return on domestic capital and so would be inefficient. This is sometimes referred to as a ‘separation’ result, since it implies that the domestic capital stock, and consequently wages and non-resource national income, are completely unaffected by the accumulation of assets from resource rents.

Although this is the benchmark, there are good reasons for departing from it.

First, a period of high resource rents, such as Chile has experienced over the past decade, may bring its own financing needs. For example, infrastructure projects may need to be brought forward, with a demand on public funds. Structural change into non-tradable goods and away

from non-resource tradables will create demand for investment and new capital spending. Some part of savings will, and should, be directed to the domestic economy. Specifically, a likely consequence of a resource boom is that the return on investment rises, temporarily raising it above the rate of time preference. High resource rents signal 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. In summary, a high income economy with perfect access to international capital markets before and after a resource windfall, should invest in foreign assets only once any domestic capital requirements associated with the discovery and associated structural change have been met.

Second, the PIH benchmark assumes that the social rate of return on domestic capital is equal to the private return on foreign assets. If this is not the case, then investment should be directed at assets that yield the social higher return. There are several reasons for thinking that in a developing economy these are likely to be domestic assets. There may be capital market imperfections. For the country as a whole, this might apply at the level of international capital markets, and for domestic firms it might arise because of poorly developed domestic financial systems. For example, the interest rate which the government needs to pay to attract foreign financing of infrastructure may include a substantial premium for the risk that the government will renege on agreed terms, but this is not a risk that the government faces if it finances the investment itself. In either case, it creates a case for directing investment to the domestic economy rather than using it to accumulate foreign assets.

Thirdly, many economies face critical shortages of economic and social infrastructure. This is because of constraints on the availability of public funds, often due to weakly developed tax systems. Even those public investments with high social rates of return may pose insuperable problems of private financing, because only a modest proportion of returns are appropriated through the tax system. Debt accumulated to finance such assets would become unserviceable. In these circumstances it can be sensible for the government to use the long term savings from resource rents to finance such assets because in its asset choice, as opposed to its debt strategy, it

can afford to be guided by the social rate of return irrespective of whether it is able to appropriate that return.

Fourthly, it is quite generally the case that the government should be looking at the social rate of return on domestic investments, i.e. the internal rate of return derived from a full cost-benefit analysis of its spending. Domestic investment may yield wider benefits due to the wedge between private and social returns created by taxes, once again suggesting the possibility of high returns.¹

These arguments all suggest that domestic investment may be a better use of funds. If the rate of return on these assets is high, but likely to fall as the capital stock is built up, resource revenues should be used to bring forward the development of the domestic economy. Norway is often taken as an exemplar of the foreign asset strategy, but foreign accumulation only began once Norway already had more invested capital per member of the labor force than any other country. Chile is not yet in that position.

For a country at Chile's level of development there is a reasonable presumption that the potential rate of return on domestic investments is higher than the modest global returns on financial assets. However, implementing efficient investment may be difficult. Countries seeking to scale up domestic spending, particularly investment, are frequently constrained by a variety of bottlenecks. A pipeline of good investment projects might be absent, and there may be a lack of capacity to design and develop projects. Project selection and cost-benefit processes may be weak, and so too the ability to procure, implement, and monitor projects. Even if projects are undertaken there may be supply bottlenecks, so that spending raises prices and buys little capital investment. This is likely to be particularly true for 'home-grown' capital: while equipment can usually be imported, structures and human capital requires domestic capacity (e.g. in the construction and training professions), all of which take time to develop.

These absorptive capacity constraints bear on the balance between investment in the domestic economy and foreign asset accumulation. If the capacity to invest well is initially limited relative

¹ This is a point emphasized by Nobel Laureate Michael Spence.

to the savings generated from offsetting depletion, then capacity needs to be built. The process of building the capacity to invest can be thought of as ‘investing-in-investing’. It has three potential components. The most evident is the capacity of government to make public investments. The IMF has a useful new measure of this capacity, the Public Investment Management Index, but unfortunately Chile is not yet included in it. Further, since public and private investment are complements, policies which facilitate private investment become more valuable. Finally, the efficiency of investment depends upon the unit cost of capital. There is a remarkably wide international variation in the cost of capital goods, both equipment and structures, and these are partly the result of policy differences. For example, in a small economy competition in the supply of equipment can be increased by integration with larger economies.

A conscious strategy of ‘investing-in-investing’ can gradually increase the capacity of the economy to absorb savings domestically. However, the level of the rents to be saved will periodically exceed the capacity to absorb them. For example, a quantum expansion in extraction due to a new mine may create a step increase in rents. In this case there is a strong case for parking savings abroad until absorptive capacity has increased, or the flow of savings diminished. Such savings are conceptually distinct from a smoothing fund: they are intended permanently to augment the domestic capital stock, but are temporarily held as foreign assets until the capacity to absorb investment exceeds the flow of savings.

In effect, the policy of ‘investing-in-investing’ endogenizes absorptive capacity. Directly, it endogenizes the capacity to absorb investment, but ultimately, once this capacity is deployed to increase investment, it endogenizes the capacity to absorb aggregate demand. Once the capacity to invest has been enhanced, the domestic rate of investment can be increased without reducing the return upon it, and the extra supply which this new capital produces enables aggregate demand to be increased without significant pressure on relative prices. It thereby provides a solution to ‘Dutch disease’ which is superior to that of the Norwegian model of saving revenues in foreign assets. The Norwegian approach is only appropriate once the domestic capital stock has accumulated to the point at which, at globally efficient standards of production, the social rate of return on domestic capital has fallen to the private return on international assets. Norway satisfies this condition, as do the petro-economies of the Gulf, but most other economies do not.

Policy Clock 3: Smoothing Expenditures

The above analysis has focused on offsetting depletion. Its key distinction was between consumption and savings. Budgets, however, work with different concepts, namely expenditure and revenue. Revenues are the sum of consumption and savings, but expenditures are the sum of consumption and domestic investment. Because it is costly to deviate from planned expenditure, savings should accommodate deviations between it and actual revenue.

That the path of future revenues is not known with certainty raises three distinct issues. One of these, uncertainty about the average rate of change of revenues, can be incorporated by an increase in the discount rate, and this in turn raises the appropriate savings rate. In effect, this is what I have already discussed in terms of the prudent response to the risk of future rent loss. 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.

Here I focus on two other types of uncertainty. One is the *intra-year* 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 other 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.

Short-Term Uncertainty and Annual Budgeting

Over a horizon of twelve months the path of physical extraction is largely known. 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, well below the mean of market expectations. Indeed, the floor price does not normally need to be 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. Further, to a modest degree it should be possible to scale back intra-year spending relative to budgeted plans at little cost: what matters is the price at which *costly* budget cuts would become necessary.

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 can 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 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 expenditure then either actual spending falls short of plans, 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, 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 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 years of high 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. One approach 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 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.

To summarize, planned expenditure is derived from the path of expected revenues. Having determined planned 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 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 plans are based on realism tempered by risk. The override requires annual spending to be reduced below its 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.

International lessons for Chile

Chile leads the world in its approach to smoothing expenditures in the face of revenue shocks. Evidently, the challenge for policy is to maintain the policy rule in the face of periods when it would be politically convenient to break it. For example, in 2011 at the onset of its oil income, Ghana established a smoothing fund. By 2014 there was already political pressure to draw it down. There had been an explosion in public recurrent expenditure, consequent on a 50 percent increase in public sector wages, and faced with a 10 percent budget deficit and a rapid depreciation in the currency, drawdown was politically attractive relative to the alternatives. Clearly, however, the fund was not established in order to prolong periods of excessive exuberance in public spending, but rather to cushion periods of revenue shortfall due to a decline in oil prices.

The best defenses against the subversion of a stabilization fund are policy inertia and a critical mass of citizens who understand the rationale for the fund. Policy inertia builds with the time that a policy has been in place. Hence, changes in the stabilization rule, even when well-justified as a genuine improvement, are likely to incur a hidden cost. In contrast, a critical mass of citizen understanding is something that has to be built and renewed. It requires an active and astute policy. The stabilization fund in Ghana is fragile not only because it is new, but because it is not understood.

The crash in world copper prices in 2009 was a valuable lesson for Chilean citizens in the benefits of stabilization. This knowledge is, however, a wasting asset unless it is regularly replenished by a running public narrative of prudence. The contrast between the persistent aversion of German citizens to inflation, and the repeated macroeconomic mismanagement licensed by the electorate of Argentina demonstrates that societies cannot rely upon automatic processes of learning such as Bayesian updating from objective experience. Most people understand causal structures neither from direct experience nor from analytics, but from well-presented narratives. One task of a central bank, however distasteful, is to master the art of such mass communication.

Implementing Policy Clocks: from Principles to Budget Rules

Budgets are 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. 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, a significant

component of revenue is from rents which are *unsustainable*. We should not expect this distinctive feature 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, they lack any mechanism for inter-temporal commitment, both for managing volatility, and for the longer horizon required for offsetting 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 IMF practice of modifying the balanced budget rule so as to exclude resource revenues – through concepts such as ‘the non-oil fiscal balance’ – has little 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.

The PIH dictates that the (constant) increment to consumption should be sustainable in perpetuity, funded from a fraction of current resource revenues and from income on the accumulated asset. This is sometimes interpreted as sustainability of the non-resource balance (NRB) but consumption and the non-resource balance are the same only if it is optimal not to use any of the revenue for domestic investment.² In general, as discussed, this will not be the case. In all but very special cases it will be optimal to have a period in which there is a relatively high level of domestic investment. This will create an NRB which appears unsustainable, implying that the NRB is an inappropriately restrictive anchor.

A rule for rising savings from resource rents

Not only should a proportion of resource rents be devoted to long term asset accumulation, but that proportion should rise over time. 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 derivative – the fact that it is unlikely to be linear – is probably

² Baumgard et al. 2012

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. 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.

Governments sometimes create specific funds into which resources for future generations are placed. When designed as Sovereign Wealth Funds composed entirely of foreign assets they breach the principle that much of the asset accumulation to offset depletion should be domestic. But even when redesigned to include domestic assets they may add to citizen confusion rather than reduce it. For example, Ghana created a Future Generations Fund into which the government paid a few hundred million dollars. However, at the same time it borrowed \$2bn through issuing sovereign bonds and reduced public investment relative to GDP. The existence of the fund created the illusion that oil revenues were being saved and thereby disguised the reality that the government response to the onset of oil was to reduce saving, using the oil as implicit collateral to borrow for consumption. It is important that citizens know the size of resource rents, and how they are used between assets and consumption. But the use of resource rents can only be ascertained by assessing the fiscal stance relative to a reasonable counterfactual estimating what the fiscal stance would be in the absence of resource rents. This is not more difficult than other fiscal counterfactuals many of which are common in policy discourse.

Rules for Managing Volatility

Above I set out the principles for smoothing expenditures in the face of revenue volatility. I 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. I 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 further principle was the need for a second line of defence: 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. Evidently, 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 money parked awaiting domestic investment to offset depletion. There is a good case for holding these assets in separate funds, with distinct rules because a parking fund needs a rule contrary to that of a smoothing fund, namely that *assets cannot be liquidated to finance consumption*.

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.

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.

Conclusion

Due to high commodity prices resource rents have become important for many countries. The management of these rents poses distinctive policy challenges that have not been faced by the major OECD countries. As a result, the standard rules of economic management have not been designed with resource rents in mind. There are no suitable practical models to follow. The handful of resource-rich high-income countries, such as the Gulf States, Norway and Australia, are structurally so different from each other, and have such radically different approaches, that none can be taken as models. Hence, the governments of resource-rich countries must think these issues through themselves.

In this paper I have focused on one distinctive and central policy issue, how much of the revenues from natural resources should be saved? I have combined the two core features which make resource revenues distinctive: they are depleting and they are volatile. An important 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, converging countries face are sufficiently distinctive that purpose-designed fiscal rules are necessary.

References:

Baumgard, T, Villafuerte M, Poplawski-Ribeiro M, Richmond C, 2012 Fiscal Frameworks for Natural Resource Rich Developing Countries, IMF Washington DC

Hartwick, John M. 1977. Intergenerational Equity and the Investment of Rents from Exhaustible Resources. *American Economic Review*, 67: 972-74.

van der Ploeg F, Venables AJ. 2011. Harnessing windfall revenues: optimal policies for resource rich developing economies. *Economic Journal*, 121:1–31