

Ugandan oil – a blessing or a curse?

John Hassler,* Per Krusell,† Abdulaziz Shifa,‡ Daniel Spiro.§

March 31, 2013

1 Introduction

This paper addresses a number of macroeconomic concerns facing a developing country that falls upon a large stock of a valuable natural resource. Our focus is on Uganda, on whose territory substantial amounts of oil were recently discovered. Although this rather unexpected event is likely to affect the lives of many—in fact, arguably most—Ugandans, history is full of examples illustrating that natural-resource discoveries cannot simply be considered manna from heaven. The finding of oil may indeed be a blessing but it can also easily turn into a curse. The purpose of this report is to analyze the key implications of an increase in revenue and to discuss some central related policy issues, such as how fast to extract the oil, how to manage the revenues, and how and when to use the revenues. Our treatment does not aim to be definitive but is rather designed to build a framework with the perspective of which one can discern and analyze the key issues.

We will start by briefly describing Uganda’s economy and some recent macroeconomic trends. We then report estimates we collected concerning the size and value of the oil discovery. In order to analyze how the oil use ought to influence the macroeconomic growth process, we then construct a tractable macroeconomic model and use it to look at the key macroeconomic tradeoffs facing Uganda. In particular, the model is well suited for a quantitative study of the tradeoffs between investing and consuming that are affected by the oil discovery.

*IIES, Stockholm University.

†IIES, Stockholm University.

‡IIES, Stockholm University.

§University of Oslo.

Since the model necessarily abstracts from a host of important issues we then move on and “think outside the theory box”, i.e., we discuss some concerns that our model cannot directly be used to analyze on a less formal level. Among these issues are risk and uncertainty, political transparency, the so-called “Dutch disease”, and whether oil revenues should be used for tax reductions. Finally, we try to draw some overall conclusions from our formal and less formal analyses.

1.1 Some key characteristics of Uganda’s economy

Uganda is a developing landlocked country in east-central Africa. The GDP in 2011 was \$16.8 billion (throughout, we use \$ to denote USD) at nominal exchange rates with a population of 35 million. This makes for a GDP per capita of \$487.¹ Using a purchasing power parity adjustment to take account of a lower domestic price level, income is almost three times as high at a GDP per capita of \$1,345. Compared to its neighbors, Uganda has a GDP per capita that is slightly lower than those in Kenya and Tanzania, but substantially higher than those in South Sudan and the Democratic Republic of Congo.

Like many other countries in Africa, growth in recent years has been higher than previously reported.² In Figure 1, we show the growth rate of purchasing power-adjusted GDP per capita in constant prices. The solid line represents the average growth rate over the preceding ten years while the dotted line represents yearly growth rates. As we see, the average growth rate has been higher over the more recent period and also substantially less volatile. The average real per capita growth rate over the last 15 years in the sample was 3.0% per year. Over the last decades, population growth rates have been fairly stable at around 3%, implying an average GDP growth rate of around 6%.

Later, we will make some comparisons to another small African country that recently started producing oil, namely Ghana.³ In Figure 2, we see that real per capita growth in Ghana shows a pattern fairly similar to that of

¹GDP levels for 2011 are taken from the World Bank at <http://http://databank.worldbank.org>

²Note, however, that the data quality is likely to be quite low. There is, in fact, suspicion that growth and GDP levels in sub-Saharan Africa are seriously underestimated; see, e.g., Young (2012). Needless to say, our quantitative analysis is conditional on the data being accurate; judging/updating the quality of the data is beyond the scope of the present study.

³See van der Ploeg, Stefanski and Wills (2011) for an analysis of Ghana and its oil resources.




Figure 1: Real per capita growth in Uganda. Source: Penn World Tables 7.1

Uganda, namely a substantial improvement since the mid-1990s: beginning at that time, the average growth rate is recorded at around 3% per year.

Government finances have been in deficits in recent years, but over the last decades, deficits have not been very large and the current (2012) general government gross debt stands at 36% of GDP. However, the current account deficits have been deteriorating at an alarming rate recently, as can be seen in Figure 3.

Let us now perform a simple decomposition of Ugandan growth per capita into the contribution of capital accumulation and productivity growth. We will use this decomposition for the calibration of the model we develop below. The idea is to use a stylized production function of the Cobb-Douglas type, i.e.,

$$y_t = z_t k_{g,t}^\gamma k_t^\alpha l_t^{1-\alpha-\gamma}, \quad (1)$$

where y_t is real GDP, z_t total factor productivity (TFP), $k_{g,t}$ the public capital stock, k_t the private capital stock, and l_t labor input, all measured in period t . The exponents γ and α are parameters assumed to be constant over time. Using data on output growth and how the capital stocks and labor evolve over time, we can back out how much output growth can be accounted for by growth in factor inputs. The remainder is then attributed to growth in productivity. This procedure would be standard if one abstracted from government capital.

The Penn World Table does not, unfortunately, contain data on capital stocks. However, it does have measures of gross investment, i.e., the sum of




Figure 2: Real per capita growth in Ghana. Source: Penn World Tables 7.1




Figure 3: Current account (solid) and General government net lending (dashed) in percent of GDP. Sources: IMF World Economic Outlook, October 2012 Online edition.

public and private capital. We can then use the relation

$$k_{t+1} = (1 - \delta) k_t + i_t (1 - \kappa_t)$$

to construct a series of capital stocks. Here δ is the average rate of depreciation, i_t is investment, and κ_t is a parameter between 0 and 1 intended to capture the assumption that some investment spending is wasted and not actually transformed into productive capital.⁴ This parameter will also play an important part in the analysis below. Finally, assuming that the growth rates of public and private capital stocks are approximately equal, we can construct the growth rate of z_t , denoted g_{z_t} , by noting that

$$g_{z_t} = g_{y_t} - (\alpha + \gamma) (g_{k_t} - g_{l_t}) \quad (2)$$

where g_{y_t} is the growth rate of real GDP per capita, g_{k_t} is the common growth rate of the two types of capital and g_{l_t} is the population growth rate. We use data from the Penn World Table for the period 1950 to 2010. We set the depreciation rate to 10% per year and κ_t to $\frac{1}{3}$.⁵ Finally, we set $\gamma = 1/6$ and $\alpha = 1/3$.

In Figure 4, we show the growth rate of total factor productivity in Uganda. As in the previous graphs, we show the average over the preceding 10 years. As we see, the growth rate of productivity increased substantially between the 1980s and the 1990s, from negative numbers to levels between 2 and 3% per year. However, there is quite a clear sign of a productivity slowdown and over the last 15 years; the average productivity growth rate has been a fairly low 0.65% per year.

Note that (2) implies that the difference between the growth rate of GDP per capita and the growth rate of productivity is the contribution from the change in capital per capita. Using this, we can conclude that productivity growth contributed to a bit more than half of the growth in per capita GDP from about the mid-1980s to the early 2000s. For example, over the period 1991–2000, the average growth in per capita GDP was 3.9 %, of which 2.4 percentage points can be accounted for by productivity growth. Over the next decade, GDP per capita grew by 2.7% per year, of which only 0.1 percentage points can be accounted for by productivity growth.

⁴We will later allow different κ s for the two types of capital.

⁵The results are quite insensitive to reasonable variations in the values of δ and κ as well as to the α 's. We also need to set the initial level of capital. However, the effect of the choice of the value for the initial capital stock vanishes fairly quickly and we report estimates of productivity only from 1975 and onwards.

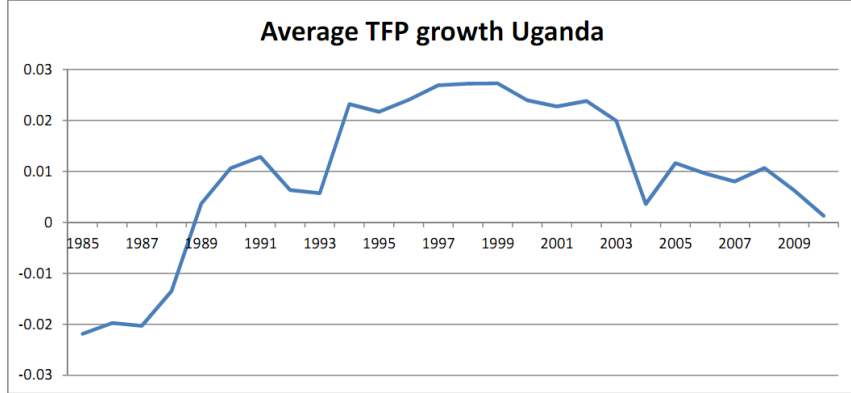


Figure 4: Ten year backward average of TFP growth Uganda.

It is interesting to note that a similar productivity slowdown did not occur in Ghana. As seen in Figure 5, the productivity growth has been fairly stable at 2% per year over the last decades.

Finally, we note that our production function in (1) allows for balanced growth—a situation where investment rates are constant and GDP per capita and the capital stocks per capita grow at the same rate. Such a balanced growth path is a key feature of the model we will present below since the economy tends to converge to such a growth path as long as the productivity growth rate and other parameters of the economy are stable. For now, we note that we can use (2) to calculate the relation between the productivity growth rate and the growth rate of GDP per capita along a balanced growth path. Using the fact that the growth rate of per capita GDP and the per-capita capital stocks are equal under balanced growth (i.e., $g_y = g_k = g_l$) in (2), we obtain

$$g_y = \frac{g_z}{1 - \alpha - \gamma}. \quad (3)$$

Under the assumption that $1 - \alpha - \gamma = \frac{1}{2}$, GDP per capita thus grows twice as fast as productivity under balanced growth. In conclusion, GDP growth was close to balanced during the 90s but unless productivity growth picks up, one might worry that growth will tend to fall in Uganda unless important structural change occurs. It is important to note in this context that the discovery of oil, however positive as a source of wealth, is in itself not source of structural change.




Figure 5: Ten year backward average of TFP growth Ghana.

1.2 Oil resources

There is substantial uncertainty about the amounts of recoverable oil in Uganda. Over one billion recoverable barrels have been discovered in the Lake Albert Rift Basin.⁶ There is also potential for more discoveries. According to the organization Oil in Uganda, a reasonable estimate is 2.5 billion barrels. After discussion with Tullow oil, we use as a somewhat conservative benchmark the assumption that there are 1.8 billion barrels that can be recovered.

Recovering and selling this oil is associated with costs in the form of exploration (around \$1/barrel), extraction (\$10-\$15/barrel) and transport (\$4-\$5/barrel).⁷ This sums to roughly \$15-\$20/barrel of costs. This is not meant to be an exact calculation of costs but suffices for producing an estimate of the economic magnitude of the Ugandan oil resource.

The quality of the oil is somewhat low; it is estimated to generate a price of around \$10 below the Brent crude oil index. With the recent Brent price being around 110 US\$/barrel, these figures would mean revenues of \$100/barrel and profits of around \$80/barrel. At current prices and the assumed costs, the oil resource thus amounts to \$180 billion in revenues and \$144 billion in profits. With a population of 35 million, the latter means \$4,100 per capita. Clearly, this is a very sizable amount compared to

⁶Stated on, e.g., the Web site of Tullow oil, <http://www.tulloil.com>.

⁷This assumes that an efficient infrastructure for transportation is constructed in Uganda. Transportation by truck is likely to be substantially more costly. Henstridge and Page (2012) claim that current truck transportation costs are \$23 per barrel.

per-capita income of \$487 per year.⁸ Since income is a flow, it may be more reasonable to compare the income that could be generated by the oil revenue to GDP per capita. If Uganda, as does Norway, invested all oil profits in a sovereign wealth fund and we assume a real rate of return of 4% per year, an added income flow of \$160 per year, i.e., an increase of current income by 33%, could be sustained forever.⁹ Another perspective on the value of the oil resource is arrived at by noting that with a growth rate of 3% per year, a 33% higher level of income is achieved after less than ten years. The purpose of this back-of-the-envelope calculation is to show that the value of the discovered oil is not large enough to itself lead to a revolutionary change in the living conditions of the average Ugandan. However, although the oil will not make the average Ugandan a rich oil sheik, if the oil resource is wisely spent, it certainly may make a difference.

2 Key macroeconomic trade-offs

2.1 When to consume and invest

Many macroeconomic questions arise when considering the added resources that the oil revenues generate for Uganda. A very basic question is what to use the revenues for: consumption or investment. This question has at least two separate dimensions attached to it. One aspect is the purely intertemporal consumption allocation issue: for any given amount of income, how much should be consumed in Uganda now and how much should be consumed in the future, by future generations? Another issue is that of production efficiency: the investment in a country such as Uganda also serves to build future production capacity, namely if the investment is in physical capital, private or public. This issue is intimately connected with the issue of growth and development.

In the section below where we provide a formal analysis of Uganda's growth process and how oil revenue ought to affect it, we provide a precise, and quantitative, answers to these questions. Here we will just briefly discuss the main conceptual issues. An important point to be made is that the

⁸ Another comparison is to the PPP-adjusted GDP of \$1,345. This level of GDP is computed using world market prices, i.e., not the actual valuations of the Ugandan economy. Thus, we prefer the comparison with local, market-valued GDP.

⁹ In reality, there is substantial uncertainty about the value of the Ugandan oil resource. First, no one knows exactly how much oil there is. Second, since the oil resource cannot be extracted immediately, uncertainty about future oil prices affects its value. We will return to this issue below.

two questions—that of intertemporal consumption choice and that of how much to invest at different points in time—do not need to be more than very distantly related. This is the case, namely, if the domestic economy is well connected with international markets for borrowing and lending and there are no obvious restrictions on any of these activities. The reason is that the international market could then be used in the consumption choice separately from the investment decisions: investment decisions affect output at different points in time but since saving does not have to equal investment the two decisions are really independent. There are many reasons to suggest that Uganda is not unconstrained in its access to international borrowing, thus creating a tighter link between consumption and investment decisions. Thus, oil revenues might be productively used by building up the manufacturing sector, for instance, but that likely implies that consumption has to be held back now in favor of only increasing the consumption in the future. Our formal model below makes this point very clear by considering different assumptions as regards the access to international capital markets.

By the same token, it is also clear that the timing of the oil extraction links in with the consumption-investment decision to the extent there are restrictions on capital/credit markets. For example, if the current situation is described by a need to increase consumption—with the argument of smoothing consumption across generations, relying on the assumption that Uganda’s future looks relatively good so that its future generations will be significantly richer than its current inhabitants—and consumption is constrained by a limited ability to borrow against future incomes, then rapid extraction is beneficial, and building up an oil fund that limits the possibilities to consume is not. If, on the other hand, current consumption is not constrained or one has a more pessimistic outlook for the future—so that saving is more desirable than consuming from an intergenerational perspective—then rapid extraction is not particularly beneficial and the costs of building up an oil fund are limited. This is likely the case of Norway, but it is far from a foregone conclusion that it fits the case of Uganda.

The benefits of investment, and what forms of investments are needed, need to be studied as well. This is of course a very difficult topic to the extent one wishes to give specific advice. The present report can therefore not presume to be very helpful in this regard, since much more detailed institutional and business knowledge would be needed than what has been possible to collect during the course of this study. The general principle should of course be a standard cost-benefit analysis: invest if the returns exceed the interest rate (whether an international or domestic rate). Our perspective on Uganda’s growth potentials is that investment will be increasingly pro-

ductive as we expect inefficiencies in the economy to fall over time. Our model below makes this quite clear at least in a general sense, but there is significant uncertainty as to how fast these efficiency improvements will arrive, whether they will be general or differ across sectors of the economy, and so on.

Finally, one cannot abstract completely from short-run aspects: though consumption and investment involve a tradeoff, they may also be complementary in that they both constitute the key components of aggregate demand. Thus, if output is demand-determined in the short run, it might be beneficial to spend resources on consumption and investment simply to raise capacity utilization. Here, of course, one can imagine different multipliers associated to consumption and investment but the main point is that both can be beneficial for the same reason. Business-cycle aspects are not considered in the formal model but are briefly touched on in the discussion following the formal section.

2.2 When to extract

In the calculation of the value of the oil resource that we made in the previous section it was assumed that the oil could be extracted at once. Though it is useful as a first approximation in the estimate, we must note this assumption is much too stark; in fact, it will take a long time to extract all the oil, even at the fastest pace possible. In conversations with Tullow oil, we have been shown their calculations of a reasonable extraction profile. The solid line in Figure 5 shows this profile measuring the flows in 1,000 barrels per day. We take it that Tullow has an incentive to propose a high speed of extraction; while there may be social reasons to slow down extraction, it is hard to see such motives for a private international oil company. Thus, we take the profile in figure 5 to represent the quickest possible extraction rate—an upper bound of sorts. As can be seen in the figure, the benchmark profile suggests a rather fast ramp-up to a maximum extraction phase of just above 200,000 barrels/day for a little less than a decade and then a slow decline.

Below, we will analyze the optimal extraction profile under various assumptions. Before that, however, we will calculate the present discounted value of extracting oil at a couple of different speeds. In particular, we contrast the quick extraction path with a path that sets the flow to a constant over the whole 45-year period covered by Tullow’s plan. By extracting 112,000 barrels a day, the same amount is produced over a 45-year periods as under the Tullow plan (that is, 1.84 billion barrels). This extraction profile is represented by the dotted line in Figure 6.




Figure 6: Tullow extraction profile in 1000 barrels per day.

To calculate the present discounted value of the oil flow, we need to make assumptions about future oil prices, the discount rate, and the evolution of costs. We take as a benchmark that the (real) discount rate is 4% per year. We assume that costs are constant and examine different assumptions about the path of future prices. A common assumption in theoretical growth models is the Hotelling rule. This theoretical result states that if resource owners are rational, the path of equilibrium aggregate supply of any resource in finite supply should lead to resource price (minus extraction costs) growth at the same rate as alternative investment opportunities. The logic behind this result is as straightforward as powerful. First, note that saving a unit of the resource, e.g., oil, for later extraction and sale can be viewed as an investment with a return given by the growth of the resource price. By not extracting and selling today, the current price (minus extraction costs) is foregone and instead the future price (minus extraction cost) is gained when the resource is extracted. If the marginal resource owner has access to an alternative investment opportunity, the market equilibrium must ensure that the return on saving oil (the price increase) is equal to the return on the alternative investment opportunity. Thus, one natural scenario is that the oil price grows at the rate of the alternative investment opportunity, in our case 4% per year.¹⁰

Under the assumption of oil price growth equal to the discount rate

¹⁰We assume that the marginal supplier has negligible extraction costs.

and extractions following the Tullow profile, the present discounted value of oil revenues minus costs is \$132 billion. Assuming a flat extraction profile instead, the value is \$136 billion, i.e., it is almost the same.¹¹ In fact, under the assumptions underlying the Hotelling result, the timing of extraction has no effect on the value of the oil resource. Under alternative assumptions, the value can be quite sensitive to the extraction path. To illustrate this, let us consider the value of the oil resource for the two extraction paths under alternative assumptions on the oil price development.

First, suppose that the oil price is constant in real terms. Clearly, in this case, delaying extraction is costly. But how costly? With a constant real oil price, the value of the Tullow extraction profile is \$76 billion, while the flat profile gives a value of \$58 billion. Thus, the delayed extraction profile implies a loss of \$18 billion, or almost one fourth, relative to the Tullow profile. We should also note that the Tullow profile has a much lower value than the (unrealistic) immediate full extraction hypothesis. This suggests that under the assumption that oil prices are constant, any delay is rather costly. In fact, the cost of delaying the whole extraction profile by one year is equal to the discount rate times the total value when prices are assumed to be constant. We conclude this analysis by noting that a delayed extraction profile makes the value of the oil resource more sensitive to variations in future oil prices.

A second consideration is that under the realistic assumption that Uganda does not have access to a perfect capital market, domestic discount rates may differ from their world counterparts. Specifically, if Uganda is credit constrained, the implicit discount rate is higher than the 4% assumed to apply on the world market. In such a situation, there is a cost of waiting to extract also if oil prices grow at the world market interest rate. As we will see below, the credit-constrained scenario seems quite realistic.

3 A macroeconomic model

In this section we provide a formal analysis, using a calibrated theoretical model, of the role of oil in Uganda's economy. The use of formal methods is beneficial in that these force discipline, both logical and quantitative. We take a growth perspective, since the main issue here is to examine the role of oil over time; thus, the maintained time horizon is rather long—on the

¹¹The fact that the flat profile gives a slightly higher value is due to the presence of extraction costs: since those are assumed not to grow, there is a (small) value of postponing extraction.

order of 100 years. With this long-run perspective we take a stand on the development process and capture the sources of long-run growth and catchup via gradual, but rather slow, elimination of inefficiencies. This approach is consistent with the recent growth literature, as for example explicated and summarized very nicely in Jones (textbook). The removal of inefficiencies implies direct benefits to output but also indirect benefits through induced capital accumulation. The model thus displays this process and how oil revenues accruing at different points in time will affect it.

We abstract from a number of important issues, but these issues are then discussed in some detail later in the text. For example, although central for the practical discussion of oil management, business-cycle fluctuations are not studied in the formal analysis but the discussion in Section 4.4 below discusses their relevance. Further, the often-discussed Dutch disease problems arising from income shocks—particularly deriving from natural resources—are absent formally in the model in this section but they can be easily introduced, at least in a basic form, and Section 4.2 discusses the associated results and interpretations.

3.1 Model description and discussion

We use an extension of the basic Cass-Koopmans model of optimal growth, the cornerstone of growth theory. This model focuses on a macroeconomic aggregate output which can be used for investment or consumption. When oil is included in the analysis, extracted oil is also assumed to be part of the aggregate output. The focus on one aggregate good implies that we abstract from international trade in different types of goods, an abstraction which is arguably not a severe one for Uganda. A version of the model—though not the baseline—describes Uganda as an open economy with access to borrowing and lending on international credit markets. In all versions of the model one presumes that oil is traded internationally at competitive prices.

3.1.1 A benchmark setup

The resource constraint is

$$c_t + i_t + i_{gt} = y_t + p_t o_t,$$

where international trade is abstracted from and where c is consumption (government and private), i is private investment, i_g is government investment in infrastructure, and y is output from domestic production (manufacturing, services, and agriculture together in one macroeconomic aggregate).

The term $p_t o_t$ refers to oil revenue, using a world price p_t and a quantity of barrels o_t . We thus distinguish the private from the public capital stock; the former is structures and equipment used in the private sector and the latter is various forms of infrastructure. Separately including government capital is important in an economy such as Uganda's; the infrastructure necessary for private production is not in place other than partially, unlike in most developed economies where the main issue is infrastructure maintenance and improvements, not basic build-up. The associated capital stocks evolve according to

$$k_{t+1} = (1 - \delta)k_t + (1 - \kappa_t)i_t$$

and

$$k_{g,t+1} = (1 - \delta_g)k_{gt} + (1 - \kappa_{gt})i_{gt},$$

where δ and δ_g are the associated depreciation rates. The parameters κ and κ_g are time-varying measures of inefficiencies in the private and public investment sectors, respectively. We thus take the perspective here that when one unit of resources is invested, a fraction disappears. The gradual elimination of these wedges will be a source of long-run growth in this economy as any unit invested will result in a higher eventual capital stock (and thus higher output) and also in higher returns from accumulating capital, potentially inducing more of this activity.

Output is produced according to

$$z_t^{1-\alpha-\gamma} k_t^\alpha k_{gt}^\gamma.$$

We thus assume a Cobb-Douglas production, giving constant shares to capital and labor income.¹² The labor input is suppressed (i.e., it can be viewed as set to 1 and to be inelastic). Population growth is not modeled here and although growth in the labor force clearly will be a source of output growth, growth in output per capita is much less dependent on population growth so we abstract from it here. Government capital, as modeled, displays complementarity with private capital: with a higher stock of it, the returns from investing in private capital go up. Importantly, z_t is a productivity parameter which both captures technical progress (making z high) and inefficiencies (making z low). Thus, temporarily low z s is thought of as

¹²We change notation slightly relative to the previous section by defining total factor productivity as $z_t^{1-\alpha-\gamma}$. This change of variables implies that the growth rate of GDP is equal to that of z in balanced growth rather than $(1 - a - \gamma)^{-1}$ times the growth of z as in (3). This change is for convenience only and does not affect any part of the analysis.

a source of underdevelopment and as the z s grow and catch up with those of the developed world, the country's economy catches up more generally in terms of capital accumulation and output since higher z s induce capital accumulation. Thus, as with the lowering of the κ s, improved efficiency has both direct and indirect benefits. Notice, finally, that we abstract from oil as an input here, since oil in this model is not of primary importance as a production input but as a source of extra revenue. This is in line with the approach taken in the literature on economic growth.

In order to evaluate welfare for this economy, we assume a representative agent with preferences

$$\sum_{t=0} \beta^t u(c_t)$$

where $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$. Here, $\beta < 1$ represents discounting, which can be interpreted both as a weight on a given individual's own future utility flows and as a weight on the utility flows of future generations. The parameter σ is a measure of how costly fluctuations in consumption are perceived to be—the higher is σ , the more painful are consumption adjustments. To find an optimal path we thus maximize the above utility function under the restrictions implied by the macroeconomic resource constraint and the capital accumulation equations.

We can implicitly define a government budget in the model, with i_{gt} (possibly plus a public part of consumption) defining spending and $p_t o_t$ defining revenues; the difference is made up by a lump-sum tax or transfer between the government and the private sector. Our focus here, however, will be that of the aggregate economy and not the government-private breakdown, since our aim is to take the perspective of the average Ugandan citizen.

3.1.2 Analysis

In the long run, the economy (i.e., its consumption, output, investments, and capital stocks) will grow at a constant rate: the rate at which z grows. Thus, in what follows we use the methods common to all undergraduate macroeconomic textbooks in the context of Solow's growth model with technological change. In this context, however, there will be a transition, as z will not necessarily grow at a constant rate at all times—it will grow fast initially, representing catch-up, and then slow down to the rate at which total-factor productivity grows in the rest of the world—and as κ and κ_g will be assumed to converge to zero. Because there will be sustained growth, we will analyze a transformed version of the model with stationary variables. For a generic variable x , we thus let $\hat{x}_t \equiv x_t/z_t$.

Having transformed all equations and the expression for the objective to be maximized, we obtain the following planning problem:

$$\max_{\{\hat{c}_t, \hat{k}_{t+1}, \hat{k}_{g,t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} (\beta^t z_t^{1-\sigma}) \frac{\hat{c}_t^{1-\sigma} - z_t^{\sigma-1}}{1-\sigma}$$

subject to

$$\hat{c}_t + \frac{\hat{k}_{t+1}}{1-\kappa_t} \frac{z_{t+1}}{z_t} + \frac{\hat{k}_{g,t+1}}{1-\kappa_{gt}} \frac{z_{t+1}}{z_t} = \hat{k}_t^\alpha \hat{k}_{gt}^\gamma + p_t \hat{o}_t + \frac{1-\delta}{1-\kappa_t} \hat{k}_t + \frac{1-\delta_g}{1-\kappa_{gt}} \hat{k}_{gt}.$$

Conveniently, thus, we can think of a planner choosing the two capital stocks over time (with an implied consumption path). Notice that we have not considered $\{o_t\}_{t=0}^{\infty}$ a choice here; we will, rather, consider different stylized extraction paths and their consequences for welfare.

The implications of the above maximization problem can be summarized in two equations:

$$\begin{aligned} \left(1 - \delta + (1 - \kappa_{t+1}) \alpha \hat{k}_{t+1}^{\alpha-1} \hat{k}_{g,t+1}^\gamma\right) \frac{1 - \kappa_t}{1 - \kappa_{t+1}} = \\ \left(1 - \delta_g + (1 - \kappa_{g,t+1}) \gamma \hat{k}_{t+1}^\alpha \hat{k}_{g,t+1}^{\gamma-1}\right) \frac{1 - \kappa_{gt}}{1 - \kappa_{gt+1}} \end{aligned}$$

and

$$\beta \left(1 - \delta + (1 - \kappa_{t+1}) \alpha \hat{k}_{t+1}^{\alpha-1} \hat{k}_{g,t+1}^\gamma\right) \frac{1 - \kappa_t}{1 - \kappa_{t+1}} = \left(\frac{\hat{c}_{t+1}}{\hat{c}_t} \frac{z_{t+1}}{z_t}\right)^\sigma.$$

The first of these equations equates the return from investments in the private sector to that in the public sector (taking into account the inefficiencies in each). Thus, this represents optimal investment behavior in the portfolio (private versus government) sense. The second equation is the usual consumption Euler equation, setting the marginal rate of transformation of resources over time equal to the marginal rate of intertemporal substitution, thus determining the solution to the optimal consumption-investment trade-offs. Eliminating \hat{c} using the resource constraint, one can combine these two equations into one second-order difference equation in \hat{k} .

We solve the model, as described by these equations, numerically for a long time horizon.¹³ The model is such that the transformed variables

¹³The solution technique relies on global, nonlinear solution and is not discussed here. The programs are available upon request.

converge to a steady state (and the untransformed variables, hence, converge to a balanced growth path).

The steady state is thus characterized by the portfolio equation

$$1 - \delta + (1 - \kappa)\alpha\hat{k}^{\alpha-1}\hat{k}_g^\gamma = 1 - \delta_g + (1 - \kappa_g)\gamma\hat{k}^\alpha\hat{k}_g^{\gamma-1}$$

and by the stationary version of the Euler equation

$$\beta \left(1 - \delta + (1 - \kappa)\alpha\hat{k}^{\alpha-1}\hat{k}_g^\gamma \right) = (1 + g)^\sigma,$$

where g is the net long-run growth in z .

3.2 The open economy

The economy that is open to international borrowing and lending has a resource constraint that reads (in transformed form)

$$\begin{aligned} \hat{c}_t + \hat{a}_{t+1} \frac{z_{t+1}}{z_t} + \frac{\hat{k}_{t+1}}{1 - \kappa_t} \frac{z_{t+1}}{z_t} + \frac{\hat{k}_{g,t+1}}{1 - \kappa_{gt}} \frac{z_{t+1}}{z_t} = \\ \hat{k}_t^\alpha \hat{k}_{gt}^\gamma + p_t \hat{o}_t + \frac{1 - \delta}{1 - \kappa_t} \hat{k}_t + \hat{a}_t (1 + r) + \frac{1 - \delta_g}{1 - \kappa_{gt}} \hat{k}_{gt}. \end{aligned}$$

where \hat{a} denotes transformed international lending (i.e., the gross level a divided by z) and r the net international real interest rate. The planning problem is thus to maximize utility subject to this constraint, with an additional choice of $\{\hat{a}_t\}_{t=0}^\infty$. This problem is easier to solve than that under the closed-economy assumption, since it delivers the following conditions:

$$\begin{aligned} \left(1 - \delta + (1 - \kappa_{t+1})\alpha\hat{k}_{t+1}^{\alpha-1}\hat{k}_{g,t+1}^\gamma \right) \frac{1 - \kappa_t}{1 - \kappa_{t+1}} = \\ \left(1 - \delta_g + (1 - \kappa_{g,t+1})\gamma\hat{k}_{t+1}^\alpha\hat{k}_{g,t+1}^{\gamma-1} \right) \frac{1 - \kappa_{gt}}{1 - \kappa_{gt+1}} = 1 + r \end{aligned}$$

and

$$\beta (1 + r) = \left(\frac{\hat{c}_{t+1}}{\hat{c}_t} \frac{z_{t+1}}{z_t} \right)^\sigma.$$

The first two of these equations can be solved directly for \hat{k}_{t+1} and $\hat{k}_{g,t+1}$, independently of the rest of the capital sequences: given that the economy can borrow and lend at r , it is optimal to simply invest in the two kinds of capital until their net returns equal r . In particular, the path of oil extraction will not influence these choices. The last equation determines

the growth rate of consumption, independently of the investment decisions. The level of consumption does depend on the investment decisions but is solved for easily in a permanent-income fashion by defining a consolidated, time-zero budget constraint

$$\sum_{t=0}^{\infty} \frac{z_t}{z_0(1+r)^t} \left(\hat{c}_t + \frac{\frac{z_{t+1}}{z_t} \hat{k}_{t+1} - (1-\delta) \hat{k}_t}{1-\kappa_t} + \frac{\frac{z_{t+1}}{z_t} \hat{k}_{g,t+1} - (1-\delta) \hat{k}_{gt}}{1-\kappa_{gt}} \right) = \sum_{t=0}^{\infty} \frac{z_t}{z_0(1+r)^t} (\hat{y}_t^* + p_t \hat{o}_t) + \hat{a}_0(1+r),$$

where y_t^* defines output given optimal investments in the two capital stocks (notice that the gross return between time 0 and time t equals $z_t(1+r)^t/z_0$ in the transformed economy). The part pertaining to consumption on the left-hand side can be simplified, given the consumption Euler equation, to equal $\frac{\hat{c}_0}{1-(\beta(1+r)^{1-\sigma})^{\frac{1}{\sigma}}}$. Thus, consumption is very easy to solve for given that the capital accumulation path is solved for from the above equations. One can, finally, obtain the evolution of net debt from the budget constraint.

3.3 Calibration

We calibrate the parts of our baseline economy not having to do with oil as follows

- $\alpha = 1/3$, representing a typical share of private capital of 1/3, with the labor share earning 2/3.
- $\gamma = 1/6$, representing a cost share of government capital of 1/6, though this capital is not traded and hence should be viewed as an externality from the perspective of the private sector.
- $\delta = 0.1$, depicting a wear-and-tear of structures and equipment of an average of 10% per year.
- $\delta_g = 0.04$, capturing a significantly lower depreciation rate on infrastructure than on private capital.
- $\frac{z_{t+1}}{z_t} = 1+g+g_0\rho_1^t$, with $g = 0.025$, $g_0 = 0.015$, and $\rho_1 = 0.85$. Thus, we take initial growth situation to be one where total-factor productivity grows relatively fast, indicating catch-up (relative to the rest of the world not explicitly modeled here). The long-run growth rate of z , and thus output, implied by these assumption is 2.5%, which is still a

bit high relative to the average rate of developed countries but the idea here is that the catchup will continue beyond the horizon considered here and this process is reasonably approximated by assuming a long-run growth rate of 2.5%. The short-run growth rate is 4% which is in line with the recent growth history of Uganda reported above. The convergence in productivity is rather slow, with persistence parameter of 0.95, thus implying that the gap away from the long-run growth rate closes by 5% per year.

- $\kappa_t = \kappa_0 \rho_2^t$, with $\kappa_0 = 0.25$ and $\rho_2 = 0.95$. Thus, the private-sector investment inefficiencies are 25% to start with and fall slowly over time (as for total-factor productivity, the convergence rate is 5% per year).
- $\kappa_{gt} = \kappa_{g0} \rho_2^t$, with $\kappa_{g0} = 0.5$, indicating that the government sector has twice the amount of inefficiencies of the private sector. We obtained the number 50% from discussions with Ugandan officials and the estimate for the private sector is simply our best guess.
- \hat{k}_0 and \hat{k}_{g0} set at balanced-growth levels consistent with a z growth of 2.5 percentage points, $\kappa = 0.25$, and $\kappa_g = 0.5$ at all past times.
- $\sigma = 1$, representing an intermediate value of the intertemporal substitution elasticity, a value that is common in the growth literature.
- $r = 0.04$, capturing a world interest rate of 4%; notice that this should be viewed as a return on capital and as a long-run average (the current low international rates being strongly influenced by the world-wide recession).
- β such that $\beta(1+r) = (1+g)^\sigma$: the long-run discounting within the country will be consistent with balanced-growth behavior at the same interest rate as in the international economy.

Thus, in summary, we view the past as quite distorted both in the private and public investment sectors, though more in the latter. However, these distortions are assumed to asymptotically vanish. We regard past and current total-productivity growth as historically high (to capture reasonable growth experience the last twenty years) and then declining toward a rate consistent with the rest of the world. Asymptotically, the closed economy will generate the same interest rate as currently in the rest of the world (the rate 4% refers to a return to capital, which exceeds that of safe bonds).

As for oil revenues, we consider the baseline scenario as the “maximum extraction path” obtained by Tullow oil (discussed above). This path implies 32,850,000 barrels extracted in an initial year of 2018 and then a rise to a maximum of 83,950,000 in just a few years (obtained in 2021), with a maintained high level of production for nine years and then a gradual tapering down with the last oil production occurring in 2062. We value oil by taking a current world price of \$100 per barrel as baseline. We then assume a Hotelling price path, thus implying a growth of 4% per year in the oil price. A more conservative price path would, for example, maintain a constant oil price, and it is of course straightforward to alter these assumptions as one wishes. A more or less rapidly increasing price path would be interesting to consider but would likely not influence our main conclusions markedly aside from the remarks that were already made in our benchmark calculations above. We set the net revenue from oil to be barrel production times the world price per barrel minus costs estimated to be \$20 per barrel; we assume these costs to be constant over time in real terms.

A crucial feature of the calibration is the size of the oil revenue relative to (net-of-oil) GDP. In terms of the model, given any normalization of z_0 , and with an initial oil price of 100, this size is obtained by selecting the barrel unit relative to total output appropriately. Uganda’s GDP in 2011 was measured at \$16.8 billion. Tullow’s estimate for 2018 of 32,850,000 barrels produced mean a revenue of $(100 \cdot 1.04^7 - 20) \cdot 32.85 \cdot 10^6$, which equals \$3.666 billion. Relative to 2011 GDP this gives 21.8%, which is how the model oil units are thus chosen. We note here that official GDP figures and a current exchange rate are used, as opposed to a PPP-adjusted measure. We deem the PPP adjustment inappropriate for the model analysis since it involves using a relative price between traded and non-traded goods which is not in line with the domestic evaluation of these goods.¹⁴ Thus, Ugandan GDP including oil would rise at the order of magnitude of 50% at the peak of the extraction path, whereas when using PPP-adjusted GDP the level is on the order of 15-20%.

A fund setup will also be examined as an alternative; the details are discussed below.

3.4 Results

We look at a sequence of illustrations, beginning with the case where Uganda proceeds without any oil income. This case is not interesting per se but it

¹⁴The model here does not have the distinction between traded and non-traded goods but the discussion in Section 4.2 does.

serves as a comparison and as a way of explaining the basic features of the economy. We then look at how oil revenues affect outcomes if they accrue according to the Tullow scheme—one that can be regarded as a maximally fast extraction—and the economy is closed, i.e., where any revenue is consumed or invested immediately. Next, this case is compared to one where there is unlimited access to borrowing and lending. We also look at a case where although the oil revenue is extracted according to the Tullow scheme, they are not used up as they arrive but rather put in a fund from which only a limited amount is withdrawn for consumption or investment each year. Although our analysis could be amended to allow any intermediate schemes for extraction and use of the oil resources, we believe these different cases bracket most of the remaining possibilities because they emphasize the possible advantages and disadvantages of different setups. Throughout we use the calibration detailed in the previous section; any additional assumptions (such as the details of the fund construction) will be highlighted below.

3.4.1 The case without oil

Under the assumption that Uganda received no oil revenue, the economy should be expected to converge rather smoothly to the balanced growth path, as TFP growth is initially high and is gradually slowing down and as the investment wedges gradually disappear. Indeed the private capital stock moves from an initial value (relative to the technology trend) of 2.1 to one of 5.4, i.e., almost a tripling, and the public capital stock moves from 1.1 to 4.7, i.e., almost a fivefold increase. These changes occur slowly both because the inefficiencies disappear slowly and because transition dynamics are a little less fast in this model, where the decreasing returns to capital broadly defined (private plus government) are weaker than in the standard model.

Let us now look in some more detail at the two key variables in their transformed versions, as graphed in Figure 7 below.

We see that the transition path is somewhat nonsmooth at the outset, reflecting an initial boost to private investment at the expense of public investment due to the fact that the falls in the wedges are larger in percentage terms for private investment. The reason is that these wedges do not accrue to the undepreciated part of capital, which is much larger for public capital, so the initial drops in the wedges work like a boost to private capital in relative terms. Thus, the initial opposite reactions of the capital stocks are somewhat surprising but logical given our setting. However, they are quantitatively unimportant in the comparison of the different setups for oil




Figure 7: Transformed capital stocks, no-oil case.

revenues that we look at below.

It is perhaps more informative to consider the (logarithms of the) non-transformed variables; they are plotted in Figure 8.

We see that all variables roughly grow in parallel, with faster transition rates for the capital stocks. The government capital stock grows the fastest since the inefficiency wedge that is eliminated for this variable is larger than for private-sector capital. The figure also illustrates the absence of business cycles: the future looks entirely smooth, of course only reflecting the assumption here that there are no shocks to the economy, foreign or domestic. Cycles are discussed in Section 4.4 below. The previously noted initial opposite movements of the two capital stocks are visible in this graph too but clearly of minor importance and consumption as well as output display no non-monotonicities.

The above paths reflect optimal transition dynamics for a closed economy. If the economy is open and can borrow and lend freely at some international interest rate (in the calibration assumed to be 4%, as discussed above), the dynamics become quite different. This case is interesting to consider not because it is the most realistic one—it is not—but because it indicates the direction in which borrowing and lending influence the economy. Figure 9 displays the results for output and consumption, in comparison with the closed-economy case.




Figure 8: Capital, output and consumption, no-oil case.




Figure 9: The closed vs. the open economies, no-oil case.

Several points are noteworthy here. First, the long-run path for output will be the same whether the economy is open or closed: the model is calibrated so that the real interest rate will be the same. We see that the open-economy output is above closed-economy output throughout (except in the first year) but with a decreasing gap. The reason why output is higher is that the internal interest rate in the closed-economy case is higher during the transition and only gradually falls to 4%, so less capital is used when foreign borrowing cannot be used to (partly) finance investment. Second, and most importantly, we see that consumption is quite a bit higher in the open-economy case initially and for several decades.¹⁵ This is because marginal utility is very high initially; Ugandan output is low at the outset compared to its future potentials, given that the inefficiencies are expected to go away over time, so it is optimal to borrow significantly right at the outset to boost consumption. The growth path for consumption in the open economy is therefore less steep than in the closed economy, and eventually the paths cross, as loans need to be paid back. Third, although welfare is not plotted, it is clear that welfare is higher in the open-economy case: consumption smoothing is beneficial, and to the extent international markets allow it and the economy can commit to paying back, borrowing is optimal. Note, however, that we are abstracting from intergenerational considerations here; in 2050, when the two consumption paths cross, many of the citizens enjoying the higher consumption with borrowing are no longer alive, leaving debt repayment to later generations. We will return to these issues below.

3.4.2 Spend-as-you-go under Tullow’s extraction scheme

We now add Tullow’s extraction scheme (refer to graph above) in the context of the economy just studied. The maintained assumption in this benchmark case is that the oil revenues (i.e., the price obtained in the world market minus the costs) go straight into the domestic budget and that these resources are allocated optimally between consumption and investment: “spend-as-you-go”. Our main focus is on a closed economy and the benchmark case is one where oil price grow according to Hotelling’s rule, i.e., at 4% per year in real terms. Such an assumption is fairly optimistic compared to many scenarios discussed in the literature. As we discussed in subsection 2.2, an alternative scenario with no real price growth would reduce the value of the oil by almost one half.¹⁶

¹⁵The initial consumption gap is as high as 71%.

¹⁶It should be pointed out that the discounting of future revenues in a closed economy does not involve a international rate but a domestic one, which we have argued is higher

Figure xyz below displays the paths for the transformed capital stocks; once again the focus on the transformed stocks is to emphasize the logic of the transition.

We see that the case with oil leads to significantly lower capital accumulation right from the outset. This may seem surprising but the reason is that in a closed economy that foresees increased income in the future there will be a consumption reaction: the “wealth effect” from the new future revenues makes it optimal to consume more now. Given an increase in consumption, if the economy is closed, there must be a fall in investment. In this case, we see that as the oil revenues come in—in 2018 according to the scheme—there is a strong rebound of investment and within a little less than ten years the capital stocks under the oil scenario are above those that would have prevailed had the oil revenue not been present. The model exaggerates the initial fall in the capital stocks somewhat since investment is assumed to be reversible and not subject to adjustment costs.¹⁷ The initial consumption boom would not be as strong if it were harder to move resources to the consumption-producing sector. Furthermore, note that we are abstracting from investments in oil extraction. Clearly, such investments are initially very large.¹⁸

The paths for (untransformed, log) consumption and output are plotted in Figure 10.

We confirm that consumption indeed is significantly higher with oil; already from the beginning of time, there is an increase in consumption of over 30%. We see that as a result, output is very sluggish—given that investment is falling so markedly—but that it catches up around 10 years after the first oil revenue is collected. Output is then permanently higher, though asymptotically the output paths with and without oil are of course identical.

As observed, the spend-as-you-go benchmark allows a rather stark increase in consumption already in the first year by reducing investment dramatically. Thus, even in a closed economy the initial consumption response is strong. Of course, the response is even stronger in one considers an open economy: there, the added oil revenue allows consumption to react

in Uganda along the transition to the balanced growth path.

¹⁷A second reason for the large initial drop in the capital stocks is that the loss experienced when investing due to κ and κ_g being positive is received back when the capital stocks are reduced. Relaxing this assumption would only affect the analysis for the initial period when capital stocks are falling, however, and would not influence the aggregate analysis more than very marginally.

¹⁸We could include these investments exogenously by subtracting them from oil revenues.

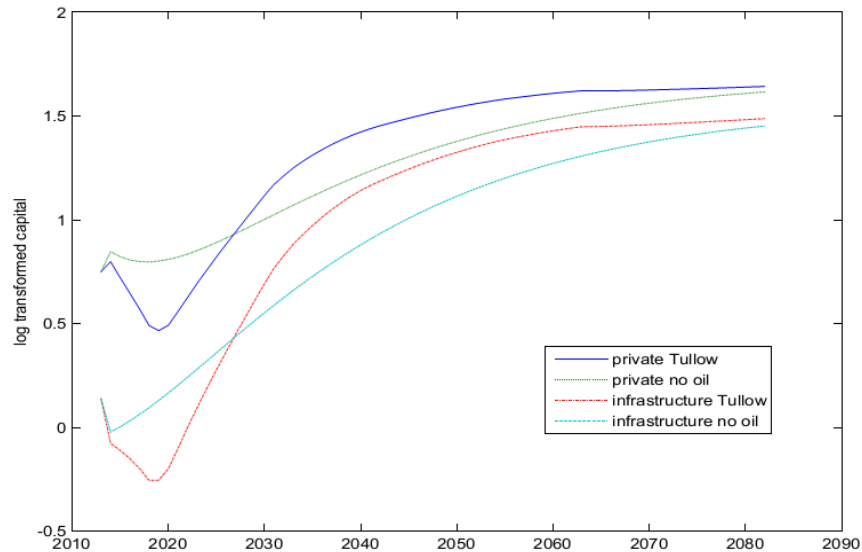


Figure 10: Transformed capital stocks, Tullow extraction, spend-as-you-go.

even more, since it is possible to borrow against all the future revenue increases. Figure 11 below compares the consumption and output paths in three economies: it displays the variables with oil and an open economy and without oil *relative to* the Tullow, closed-economy benchmark.

Looking at consumption first, clearly there is a major additional boost to consumption initially if borrowing is unrestricted.¹⁹ The open-economy boost to consumption is also prolonged and the consumption path of the benchmark closed-economy case only catches up at around 2070. One reason for this is the effect on output: output also experiences a major boost in the open economy, since it is now possible to invest immediately at a lower rate, without any adverse impact on consumption. Thus, overall welfare is increased both by allowing better consumption smoothing—in particular, consuming early when marginal utility is very high—and by increasing production in a present-value sense. The figure also graphs the no-oil cases for comparison.

¹⁹The boost is almost 50% initially.




Figure 11: Consumption and output, Tullow extraction, spend-as-you-go.

3.4.3 A fund

Finally, we look at the case where an oil fund is used. The following fund construction is adopted: (i) each year 4% of the fund is withdrawn for consumption and investment and the rest is invested in the world markets at 4%; and (ii) extraction according to the Tullow scheme is maintained and added to the fund as the revenues accrue. Thus, the fund construction allows significantly less consumption smoothing than does the benchmark scheme at the same time as it prevents output from being sluggish. Figure 12 displays the results.

Consumption is higher for the benchmark case than in the fund case until about 2045, and by a significant percentage amount (about 20% until around 2030, when the difference falls). The other side of this coin is that output does not remain stagnant but keeps growing at a decent rate.

3.5 Conclusions from the model

Our quantitative optimal growth model clearly puts the oil revenue management in a perspective: that of the overall development path of Uganda. First, and quite crucially, given that our maintained hypothesis is continued growth and catchup, with successive elimination of inefficiencies in the

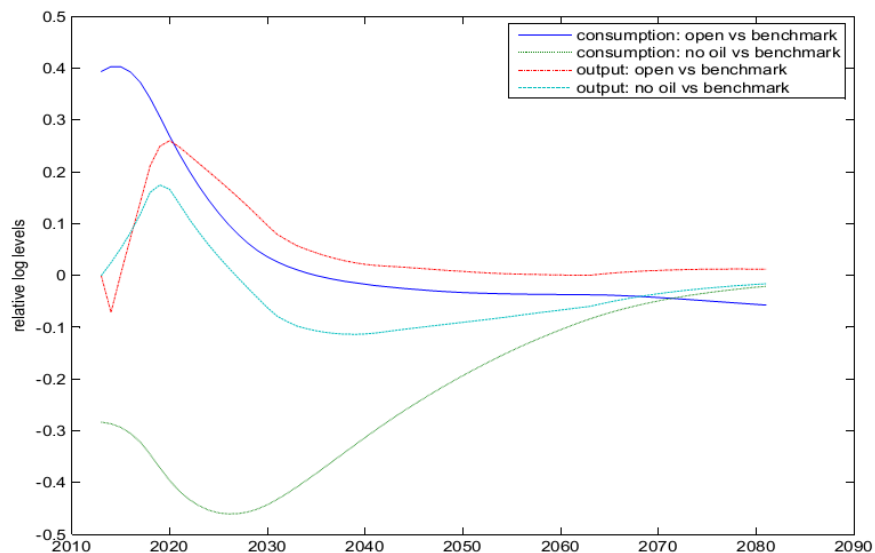


Figure 12: Relative consumption and output levels, open vs. benchmark.

production of consumption and investment (particularly for government capital), there is a strong need for consumption smoothing. That is, the marginal utility of consumption is naturally very high initially and given the future output increases any restrictions on consumption will lower welfare. In particular, if it is possible to borrow in the international markets merely to fund consumption, it would be good to do so.

Second, the implications for oil revenue management are clear, at least on a qualitative level: the oil revenues will help smooth consumption, and foreseeing their arrival, consumption will, and should, rise already today. In a closed-economy context, this means that domestic output production is hurt since it lowers investment. Moreover, comparing the case when one can borrow to that when one cannot, borrowing allows further consumption boosts early on since the economy borrows against the future oil revenues, and output will not be held back since the relevant real interest rate is the international one and not the much higher local rate.

Third, considering the case of a fund that holds back consumption it is clear that the initial boost in consumption is less marked, as is the drop in output growth. However, welfare under this scheme is lower than under the benchmark precisely because of the effects on consumption. It is also clear that a more consumption-generous fund scheme would go part of the way toward the benchmark. How large are these effects, i.e., how large is the

loss in reducing consumption smoothing under the fund scheme considered in the section above? Arguably not giant. We saw significant early losses from using the fund but, of course, later gains. Any such losses must be weighed against other drawbacks of the spend-as-you-go benchmark scheme, such as a possible Dutch disease, strains on the local consumption-producing sector, incentives for inefficient/opportunistic use of the revenues, and so on. Thus, the basic observations and recommendations we deliver here come with a number of qualifications and these are discussed in the following sections.

4 Considerations outside of the model

The model we constructed and analyzed in the previous section is obviously stylized and necessarily abstracts from a number of important considerations. In this section, we will go beyond the formal analysis and discuss a number of issues to which particular attention must be paid.

4.1 Transparency and accountability

Countries abundant in natural resources such as oil and metals have experienced highly disparate economic, political, and social developments. There is a vast amount of studies that document how countries are affected by natural resource discoveries (see, e.g., van de Ploeg, 2011, for a survey). Some experiences are clearly positive ones: the added resources lead to improvements along most economic and social dimensions. However, for other countries—unfortunately a majority—the experience is instead rather negative. The consensus today is that the largest risks arise in the political arena. The rents from the resources may, and often do, spark or reinforce corruption at various political and social levels and lead to the undermining of democratic institutions, all with the purpose of gaining control over the resource rents. These effects, in turn, trigger economic stagnation, inequality and sometimes even armed conflicts. To uphold and develop the political institutions is therefore of first-order importance also for Uganda.

One needs to be extremely wary of the possibility of these negative developments and it is important to realize that they are not set in stone: there are measures one can undertake to minimize the risks that they will surface. One problem that can be addressed has to do with transparency. One example of lack of transparency is the sharing agreements of oil revenues in Uganda: they are not public. This clearly limits the ability of media and the citizens of Uganda to scrutinize the agreements and to investigate whether

they contain problematic elements and whether they address all the important issues in a proper way. An explanation often raised to defend the lack of transparency is that the sharing agreements contain sensitive information for the oil companies. This explanation does not, however, seem satisfactory. First, our clear impression from discussing the agreements with the oil companies is that these companies themselves do not ask for confidentiality. Second, there are many countries where the sharing agreements are public without hindering international investments.

The recent oil bill that was approved by parliament gives the Minister of Energy and Mineral Development the authority to sign and revoke agreements with oil companies. To avoid the risks of considerations not in line with the welfare of the Ugandan people, i.e., an undue influence on the oil exploration and the extraction process, it would be preferred if the rules governing the formation and signing of agreements were drawn by the minister and approved by parliament but, importantly, that the execution of these rules would be carried out by an independent authority. This way (i) the authority can be held accountable if they fail to implement the rules correctly and (ii) the politicians can be held accountable if they create rules which are not publicly approved. We want to emphasize that the lack of transparency certainly is a problem also in developed countries.

A closely related problem in resource-rich countries is that large revenues may make spending decisions worse from a social point of view. Adverse partisan influence over these decisions is hard to avoid also in developed countries and call for caution in Uganda as well. First of all, oil income creates the risk of using more of the revenues for public spending in election years to boost the popularity of incumbent politicians. In Uganda, following the oil bill, the parliament decides on how much of the oil revenues each year that will be invested in the sovereign wealth fund and how much that will be spent right away. Likewise, the parliament decides on how much that will be taken out from the sovereign wealth fund. The beneficial democratic effect of this is obvious and it allows for using the sovereign wealth fund as a tool for mitigating business-cycle fluctuations, an issue we will return to below. However, it naturally also creates a temptation to over-spend in election years. To avoid this, some countries have created a spending rule which states how much can be used in a certain year. For example, the Norwegian spending rule stipulates that all of the revenues from the oil go straight to the SWF and that no more than 4% of the fund's value can be used in a given year. This precise formulation implies a spending of oil revenues which is increasing over time. As we showed in the theoretical sections, such a path may imply too low a consumption flow for current generations and should

thus not be taken as a blueprint for a developing country such as Uganda. However, a modified version of such a rule may still be useful in order to avoid political temptations. One possibility would be a rule that specified a path for investment in the sovereign wealth fund over time. Currently, such a rule or consensus does not seem to be available in Uganda but it would be desirable to arrive at one.

An issue related to how to use the resource income is the selection of specific investment projects. Earlier research has documented that resource spending quite often goes toward projects with low returns, motivated by pleasing various political groups or electorates. A key optimality condition that came out of our theoretical model was that the return on public and private investments should be equalized. Since the return on public investments are not measured by markets, they have to be calculated by economists trained at cost-benefit analysis. These skills are available in the Ugandan authorities. However, the fact that individual public investment projects may benefit some groups of individuals more than others creates strong incentives to distort these calculations. In addition, of course, outright corruption is an obvious danger.

The above considerations imply that it is necessary to strengthen the state comptroller and other authorities monitoring the conduct of government, politicians and the bureaucracy. We want to emphasize that these are measures that should be undertaken before the resource revenues start flowing, as strong institutions will not only be beneficial for avoiding corruption but also for protecting themselves against the risk of being undermined by politicians with personal agendas. Deciding on how to spend on public investments is a political task that hardly can be delegated to an independent agency, at least when the investment project is large. However, a rule stating that public investments over a certain size needs to be evaluated by an independent agency can increase transparency and enhance the quality of decisions without compromising democratic principles. Here, it may be advisable to use international consulting agencies until domestic institutions with sufficient competence and independence are built. It may also be reasonable that when large infrastructure projects are proposed, the runner-up alternative in terms of social profitability is also presented.

4.2 The Dutch disease

One often raised concern is a detrimental effect that added natural-resource revenues can have on other sectors. This mechanism is the so-called Dutch

disease.²⁰ This concept captures the tendency that oil extraction leads to a currency appreciation, lowering competitiveness in other export sectors. In an analysis of this issue, it is important to first note that the mechanisms behind a loss of competitiveness in other export sectors following an expansion of oil exporting exist also in economies that function perfectly. After a windfall income gain, domestic demand for non-tradables produced domestically increases. To meet the higher demand, output expands, which requires a resources reallocation from the export sector to the non-tradables sector. This is achieved by a real appreciation that reduces the relative profitability of non-oil exports. Thus, an increase in the real exchange rate should not necessarily be judged as a failure or an unwanted side effect. In this sense labeling it a “disease” may be quite misleading.

To illustrate this, consider a very simple static model where individuals consume two goods c and s . Good c is a tradable good and its price, normalized to unity, is given from the world market. The price of the non-tradable is denoted p and this is also a measure of the real exchange rate. A representative household derives utility from the two goods according to

$$U = \ln(s^\gamma c^{1-\gamma}).$$

Here, we can think of $C \equiv s^\gamma c^{1-\gamma}$ as a measure of aggregate consumption.

Domestic production of good c is exposed to foreign competition: it is import-competing. However, the production of good s is shielded from competition from abroad. The representative household purchases the two goods under a budget constraint

$$ps + c = w$$

where w is income. Now, it is straightforward to show that the representative household will spend an income share on each good that is determined by the parameter γ . Specifically, γw is spent on the non-tradable and $(1 - \gamma)w$ is spent on the tradable. Thus, $ps = \gamma w$ and $c = (1 - \gamma)w$, implying that

$$p = \frac{\gamma}{1 - \gamma} \frac{c}{s}. \quad (4)$$

Suppose now that economy finds oil so that w goes up. Consider two possibilities, the first being that the supply of the non-tradable is inelastic. To make the point clear, suppose it is perfectly inelastic. Then, by assumption, s cannot change and all the oil revenue is used on the tradable. In this

²⁰See, for instance, Corden & Neary (1982) and Matsuyama (1992).

case, we see from equation (4) that p , the real exchange rate, must increase in exactly the same proportion as consumption of the tradable good. We also note that an increase in w by x percent leads to an x -percent increase in the tradable but only a $(1 - \gamma)x$ -percent increase in aggregate consumption since the output of the non-tradable cannot respond to the increased demand.²¹

The second possibility is the opposite one: that resources (e.g., capital and labor) can move quite freely between the two production sectors so that shifts in relative demand is met by equal shifts in supply. In this case, the relative price of the two goods will not respond at all to shifts in relative demand, i.e., p is constant (here it is normalized to unity). Then, from (4) an increase in the oil revenues will cause both c and s to increase in the same proportion as w , i.e., by x percent. In this case, aggregate consumption increases by x percent too! Thus, welfare increases more than in the previous case, and the difference is larger the larger is γ .²² Note that in the latter case, when the welfare increase induced by the finding of oil is maximal, the transfer of production capacity from the tradable sector to the non-tradable was also maximal.

An important conclusion from this reasoning is that increased oil revenues should be allowed to lead to a transfer of factors of production from the tradables sector to the non-tradables sector. A second equally important conclusion is that this mechanism is stronger the higher is γ , the income share of non-tradables. From this it also follows that when the supply of non-tradables is less than perfectly elastic, which certainly is the realistic case, any distortionary policies that make the non-tradable sector larger than optimally also make the social value of oil revenues smaller. In other words, the social value of opening up non-traded sectors to foreign competition increases when oil is found.

Of course, this simple model abstracts from many important considerations. One such consideration is the presence of externalities in the non-oil tradable sector. An example would be a non-oil export sector where human capital accumulation or technology adoption play key roles and involve spillovers. In such a case, a reduction of activities in these sectors may

²¹Another way to say the same thing is to note that measured using world market prices, gross national income increases by x percent, but the domestic price level increases by γx percent. This occurs since the cost measured in terms of tradable goods increases by x percent on a share γ of the consumption basket.

²²We do not know the share of non-tradables in the Ugandan economy but there are reasons to believe that it is large, perhaps too large to be efficient. In fact, Muhumza (2011) argues it is as large as 75%.

hamper growth both in the short and the long run. Similarly, there may exist irreversibilities where a lowering of investments in a sector makes it difficult to re-stimulate it once the oil runs out. In itself, such irreversibilities do not invalidate the arguments put forth above. However, there may be other imperfections that lead to such irreversibilities not being properly taken into account by markets and governments. For example, it is possible that individual occupational and educational choices—perhaps because of lack of information—do not fully take into account the temporary nature of the change in relative profitability of different occupations associated with a real exchange rate appreciation due to a temporary boost in export revenues. In such a case the Dutch disease may cause problems also in the long run.

Political considerations are important in this context as well. One can easily imagine that the pressure to remove inefficient regulation in the export sector and undertake other measures to enhance its competitiveness falls as the economy finds a new source of export revenues. This is a clear danger that needs to be addressed at its root, i.e., by making sure that the reform pressure remains high. Measures focused on affecting the real exchange rate are thus not likely to be helpful. However, a strong real appreciation associated with serious problems of competitiveness may be a warning signal that productivity enhancing measures need to be strengthened. The simple back-of-the envelope calculation in section 1.2 may serve as a warning: if growth over just a few years is permanently lost due to bad economic policy, the cost of this can easily be of the same order of magnitude as the value of the Ugandan oil.

Another potential problem associated with a real exchange rate appreciation is that it can have unwanted distributional consequences. A real exchange rate appreciation by definition implies that the relative price of domestically produced goods increase. To the extent that low-income households consume a larger share of such goods than does the average household, real income inequality may rise. A real exchange rate appreciation will tend to raise the relative price of domestically produced non-traded food. It is well known that this can lead to increased social tension and even severe violence. Also, as was discussed in a previous section, research has documented the strongest resource curse effects coming through an institutional channel where resources lead to more corruption and institutional decline and an amassment of the oil profits in the political and economic elite, which exacerbates the economic inequality. These problems may certainly overshadow any macroeconomic problems associated with the Dutch disease.

4.3 Lowering taxes

With revenues from oil, it is possible to lower taxes while keeping total government revenues unchanged. Would such a policy be efficient?

To answer this question, one needs to ask how distortionary Ugandan taxes are currently. On the one hand, general government revenue is low compared to that in developed countries, totaling 15.6% of GDP in 2010.²³ This suggests that tax rates, and therefore distortions, are low. On the other hand, low total tax revenues may be the result of a combination of high tax *rates* and pervasive activities to avoid taxes. In such a situation, the tax distortions may be high despite the low revenues.

A complete description of the Ugandan tax system is beyond the scope of this report. However, we note that at least some Ugandan tax rates appear to be high also from an international perspective. According to the web site http://www.tax_rates.cc, the standard VAT rate is 18% and the corporate income tax is 30%. A clear possibility is also that actual tax payments are concentrated in certain sectors and larger firms. Similarly, taxes are naturally higher in the formal sector, which is also typically the most important for technology adoption and export activities. This creates a wedge that not only distorts the economy in a static sense but also hinders development and reduces growth. In fact, it is precisely this mechanism that is argued to be a main cause behind the low degree of competitiveness of the Greek economy. Greece has a very high share of self-employed workers, arguably due to a highly distorting effect of tax rates in combination with loopholes that more easily can be taken advantage of by small firms and the self-employed.²⁴ As in Greece, there is also reason to suspect that taxes create distortions that disproportionately affect export- and import-competing sectors, thus constituting a prime suspect behind tendencies toward real exchange rate appreciations and current account deficits.

Clearly, higher state capacity and higher-quality public institutions generally make it possible to create a more efficient tax system, i.e., a system that can generate a given revenue with less distortions. However, one may also consider the reverse causal chain. It has been emphasized in the literature that financing government expenditures with taxes rather than with foreign aid may lead to a more effective democratic control over how government resources are used.²⁵ The idea here is that citizens may have a tendency to care more about how their own tax payments are spent than

²³Source:IMF, WEO Database, October 2012.

²⁴See Corsetti et al. (2011).

²⁵See Moss et al., (2008).


about other government revenues. This creates an argument in favor of taxes and may have the implication that tax distortions are lower than without this mechanism. A tax system with broad tax bases where a large share of the population participate in the financing of public goods may, thus, generate high levels of democratic control. In contrast, a system full of loopholes is more likely to generate activity aimed at reducing tax payments.

Taxation capacity is not built over night and efforts to make the system more efficient should not be diminished. As discussed in the previous subsection, there is an obvious risk that windfall income, like that from expanding oil revenues, reduces the pressure to deal with a badly constructed tax system. It is very important that the Ugandan economy not pretend that “it can afford not to undertake reforms” Our impression that current taxes are quite distortionary and growth-hampering leads us to conclude that the reduced financial pressure on the government coming from future oil revenues should be used to reduce taxes on the formal sector of the economy. Reductions in profit taxes and capital income taxation should, therefore, be a very welcome aspect of a tax reform. By using oil revenues to reduce taxes, there is thus a possibility of a “double dividend” from the new oil revenues, since the reductions in tax rates have secondary benefits on Ugandan citizens and are likely to be growth-enhancing.

4.4 Income volatility and sovereign wealth funds

The oil price is notoriously volatile, as seen in Figure 12. Over as short a period as the last five years, the monthly average crude oil price has varied between over \$130 per barrel in July 2008 and close to \$40 per barrel during the spring of 2009. As the financial crisis has been calming down, the oil price has stabilized somewhat but certainly remains extremely volatile relative to, e.g., manufactured goods traded on the world market. The exact causes of these fluctuations are not known but both long-run and short-run expectations of supply and demand seem to play important roles. Clearly, the safest prediction about future oil prices is that they are likely to remain volatile and uncertain for the foreseeable future.

The price volatility has (at least) two important implications. The first is that the income stream from a given smooth extraction path becomes volatile. Since extraction is partly determined by physical constraints and by capital investments that are made many years in advance, it is not economically reasonable to change the extraction path in order to perfectly smooth revenues. The second implication is that the total value of not yet extracted resources is unknown. Since our theoretical model abstracted from



such uncertainty, we will address some of the concerns here.

Clearly, a government budget should not be built on using the very fluctuating flow of oil revenues hand-to-mouth. A well-functioning international capital market could in principle be used to smooth stationary short-run fluctuations in oil revenues. However, as discussed in the theoretical section, access to such a market is likely to be constrained. This creates a motive of its own to instead build a buffer stock in the form of a sovereign wealth fund which would work as a buffer between the economy and the oil revenues. Inputs into the fund may fluctuate while withdrawals from the fund can be made smoothly. Initially, however, when the sovereign wealth fund is still small, its buffering capacity is naturally limited.

There is substantial evidence that government consumption in developing countries is pro-cyclical and that this increases business cycle volatility; see Ilzetzki and Vegh (2008). There are reasons to believe that this has negative effects on growth, in particular in developing countries with less developed financial markets.²⁶ Oil revenues may be a way out of this. But because oil prices are so unstable, it requires the use of a stabilization fund. Otherwise, oil revenues may worsen the situation and be harmful for growth.

Norway has one of the world's largest oil funds. From the end of the year 2000 to the end of year 2011, the fund increased from 26% to 121% of Norwegian GDP. The current value is around \$700 billion.²⁷ The purpose of

²⁶See Aghion et al. (2009) for evidence on a negative relation between real exchange rate volatility and growth.

²⁷See Steigum (2013, forthcoming) for an analysis of the macroeconomics of sovereign wealth funds.

the Norwegian fund is “to avoid excessive and non-sustainable government spending of petroleum revenues” (Steigum, 2013). In Norway, it seems reasonable to assume that the private as well as the public sector has good access to the international capital market. Therefore, using the large oil revenues to finance domestic investments would to a large extent only crowd out other sources of financing. Therefore, quickly building a large sovereign wealth fund that invests in foreign fixed income securities and shares was a quite natural choice. According to the fiscal guidelines set by the Norwegian parliament, the government can withdraw 4% of the asset value of the fund per year to be used in the budget. An important feature of the Norwegian fiscal framework, however, is that these 4% are measured as an average over the business cycle. Thus, the fund is used as a source of stabilizing government revenues, even allowing for counter-cyclical expenditures.

The second consideration we want to bring up is the fact that an uncertain future oil price makes the value of not yet extracted oil resources unknown. There are reasons to believe that the oil price is driven by a non-stationary stochastic process. This implies that oil price fluctuations do not tend to average out in the long-run. Instead, uncertainty about the oil price should increase as the forecast horizon is extended. This may create an incentive to reduce uncertainty by extracting quickly. Probably more importantly still, this creates a precautionary motive for saving, counteracting any borrowing against expected future oil incomes. This is an important consideration that was not included in the formal analysis in section 3.

For all the above reasons, we conclude that there are important arguments in favor of implementing a Ugandan sovereign wealth fund. There is substantial academic work on sovereign wealth funds.²⁸ There is also substantial practical experience. An important example to study is Ghana’s Petroleum Funds that were set up shortly after commercial oil extraction began. Consistently with the two considerations mentioned above in this subsection, Ghana has separated its wealth holdings into two separate funds, one intended to help smooth government budgets in the short and medium run and one intended to consider long-run objectives, like inter-generational fairness. In Norway, it was decided that one wealth fund could serve both purposes by imposing the relatively flexible rule that as an average over the business cycle, 4% of financial assets could be used by the government. Such a rule allows flexibility but is problematic from an accountability point of view. Since business cycles are irregular and since it is a matter of judgement

²⁸See, for example Das et al. (2009). Bartsch (2006) focuses on the first of our considerations, namely the stabilization motive.

whether a given change in GDP is cyclical or more permanent, it is difficult to evaluate whether the 4% rule is actually followed in any given year. Due to these considerations, the Ghanaian example seems more suitable to Uganda. With a separation of the funds, limits are placed on how much resources can be used for stabilization purposes; this is likely to enhance accountability. Portfolio considerations too may speak in favor of a separation. A stabilization fund should aim for a higher level of liquidity than a fund created for long-run objectives. Therefore, the benchmarks against which the two types of funds should be measured should likely differ. A separation thus makes it easier to create clear benchmarks which can facilitate fund accountability.

Compared to Norway's case, where most of oil revenues are currently saved for future generations, 50-70% of the oil revenues in Ghana are not put in the oil fund but used immediately by the government (the "spend-as-you-go" part). Of the remaining 30-50%, about 2/3 are put in the stabilization fund.²⁹ A reason for the relatively large share put into the building of a stabilization fund is for the fund to grow large, and therefore useful, fairly quickly. Furthermore, as noted above, to the extent that a stabilization fund can stabilize the macroeconomic fluctuations and possibly be growth-enhancing as well, building such a fund may reasonably be a high priority.

5 Conclusion and recommendations

We will now draw some tentative general conclusions from our analysis. A first important conclusion comes directly from the theoretical model. The value of the oil discovered in Uganda is large. We have made relatively optimistic assumptions about the general outlook in Uganda, in particular about the future oil price and about the possibility for Uganda to catch up with the developed world in terms of TFP and investment efficiency. Under such assumptions about the future, the model delivers the conclusion that there is quite a strong argument for letting current generations share in the prosperity promised by the future oil revenues. On the other hand, we have emphasized that future income is not safe manna from heaven: it is uncertain and requires that many challenging choices be made successfully and temptations be avoided systematically. History is full of examples of resource-rich countries that have found themselves in a situation where optimistic forecasts of future wealth have been replaced by despair—gold turning into sand, or something worse than sand.

²⁹See, van der Ploeg (2011).

Our recommendation is therefore, first, to exercise caution in consuming the expected proceeds from oil revenues, more so than what our model says is optimal. The large current account deficits need to be analyzed and government deficits should not be allowed to grow too fast. It is well known, not the least from the recent developments in the Euro area, that large debts can be dangerous and often are destabilizing. Nevertheless, we maintain that substantial amounts of oil revenues should be used for the benefit of current generations. How much exactly is, of course, a political decision. However, we insist on the importance of transparency in making this political decision. One possible way to achieve this is to set aside a share of revenues in a sovereign wealth fund that invests in a broad world market portfolio as, e.g., practiced by the Norwegian oil fund. In order to have a well functioning process for making decisions about how to use oil revenues, information about contracts and forecasts must be made public in a much freer way than currently.

We have also argued that oil revenues can yield a “double dividend” by making it possible to reduce distorting taxation. We do not have sufficient knowledge about the Ugandan tax system to be able to suggest a specific tax reform, but reducing capital and corporate income tax rates may be particularly important for growth while a reduction of value added taxes may be more important for static efficiency and perhaps equality.

Another possible form of double dividend is the use of oil revenue as a stabilization device. Properly employed, they can help the governments overcome the tendency for procyclical fiscal policy that is endemic in the developing world. However, in itself, the oil revenue is likely to be both highly volatile and procyclical. Therefore, Uganda is well advised to set up an explicit stabilization fund that can be used for this purpose. The stabilization fund and the wealth fund, moreover, should be set up separately. Since they have different purposes and should be evaluated against different benchmarks, such a separation is important for ensuring accountability. Here, Ghana can provide an important example but lessons can also be drawn from as far away as Norway. In particular, Norway provides a good example when it comes to transparency and clarity of goals, the investment practices, and accountability.

A final recommendation is to introduce a law that requires cost-benefit analyses to be undertaken before large public investments are decided on. Such analyses should be made by independent agencies. In the near future, it seems reasonable to use international consulting firms for this. However, independent domestic agencies for the evaluation of how public resources are spent should also be set up. Moreover, this endeavor should not wait

until oil revenues start to arrive.

All in all, after writing this report, we are optimistic: it seems possible, and in power of Uganda's citizens, to make oil a blessing and not a curse.

6 References

Bacchetta, P., R. Rancière and K. Rogoff, (2009), "Exchange rate volatility and productivity growth: The role of financial development". *Journal of Monetary Economics*, 56(4).

Bartsch, U., "How Much is Enough? Monte Carlo Simulations of an Oil Stabilization Fund for Nigeria", IMF Working Paper 2006:142.

Corden, W. M. and Neary, J. P. (1982), "Booming Sector and De-industrialisation in a Small Open Economy", *Economic Journal* 92 (December), 825-48.

Corsetti, G., M. P. Devereux, J. Hassler, G. Saint-Paul, H.W. Sinn, J.-E. Sturm and X. Vives, (2011), "Chapter 3: Greece", EEAG Report on the European Economy 2011.

Das, U. S, Y. Lu, C. Mulder, and A. Sy, (2009), "Setting up a Sovereign Wealth Fund: Some Policy and Operational Considerations", IMF Working Paper 179:2009.

Henstridge, M. and J. Page, (2012), "Making a Modest Boom: Oil Revenues in Uganda", OxCarre Research Paper 90, Oxford Centre for the Analysis of Resource Rich Economies.

Ilzetzki, E., and C. A. Vegh, (2008), "Procyclical Fiscal Policy in Developing Countries: Truth or Fiction", NBER WP 14191.

Jones, C. and D. Vollrath (2013), *Introduction to Economic Growth*, W.W. Norton.

Matsuyama, K. (1992), "Agricultural Productivity, Comparative Advantage, and Economic Growth" *Journal of Economic Theory*, 58(2): 317-34.

Moss, T., G. Petterson and N. van de Walle, "An Aid-Institutions Paradox? A Review Essay on Aid Dependency and State Building in Sub-Saharan Africa", in ed. W. Easterly, *Reinventing Foreign Aid*, MIT Press, Cambridge USA.

Muhumuza, F. K., "Like Europe, Uganda needs major economic reform", Daily Monitor, Uganda, August 28, 2011.

Steigum, E., (2013 forthcoming), "Soveriegn wealth funds for macroeconomic purposes", Report for the Swedish Fiscal Policy Council, Stockholm, Sweden.

van der Ploeg, R. (2011), "Natural Resources: Curse or Blessing?", *Journal of Economic Literature*, Vol. 49, No. 2, pp. 366-420(55).

van der Ploeg, R., R. Stefanski and S. Wills, (2011), "Harnessing Oil Revenues in Ghana", IGC.

Young, A., (2012), "The African Growth Miracle", *Journal of Political Economy*, 120:4.