Should Developing Countries Constrain Resource-Income Spending? A Quantitative Analysis of Oil Income in Uganda

John Hassler, Per Krusell,* Abdulaziz B. Shifa,** and Daniel Spiro****

ABSTRACT

A large increase in government spending following resource discoveries often entails political risks, inefficient investments and increased volatility. Setting up a sovereign wealth fund with a clear spending constraint may decrease these risks. On the other hand, in a capital scarce developing economy with limited access to international borrowing, such a spending constraint may lower welfare by reducing domestic capital accumulation and hindering consumption increases for the currently poor. These two contradicting considerations pose a dilemma for policy makers in deciding whether to set up a sovereign wealth fund with a spending constraint. Using Uganda's recent oil discovery as a case study, this paper presents a quantitative macroeconomic analysis and examines the potential loss of constraining spending through a sovereign wealth fund with a simple spending rule. We find that the loss is relatively low and unlikely to dominate the political risks associated with increased oil spending. Thus, such a spending constraint appears well warranted.

Keywords: Economic Development, Macroeconomic Dynamics, Oil, Resource Curse, Sovereign Wealth Fund, Uganda

https://doi.org/10.5547/01956574.38.1.jhas

1. INTRODUCTION

Countries making resource discoveries often find themselves facing a number of important and difficult decisions. One such decision regards whether to set up a sovereign wealth fund—or more generally, determining a time profile for the usage of the resource income. Using Uganda's recent oil discovery as a case study, we quantitatively analyze this question in a developing country setting featuring borrowing constraints and capital scarcity.

A standard recommendation is that a capital-scarce developing country, facing a borrowing constraint, should use a large share of its resource revenue to boost current spending (see, e.g., van der Ploeg and Venables, 2011). This is of course intuitive given that both marginal utility from consumption of the currently poor generation and the return from domestic investments is relatively high. Thus, a " spend-as-you-go" scheme—where all of the oil revenue is spent as it arrives—may seem appropriate to increase spending both on current consumption and domestic capital formation.

On the other hand, a large inflow of income from natural resources often leads to corruption and various negative political effects as politicians, officials and elites try to get part of the resource

* Institute for International Economic Studies, Stockholm University.

** The Maxwell School, Syracuse University.

*** Corresponding author. Department of Economics, University of Oslo. Affiliated with Oslo Center for Research on Environmentally Friendly Energy. E-mail: daniel.spiro@econ.uio.no.

rents (van der Ploeg, 2011; Vicente, 2010; Sala-i-Martin and Subramanian, 2003; Collier and Hoeffler, 2004). These effects, in turn, trigger economic stagnation, inequality and sometimes even armed conflicts.[1] A closely related problem in resource-rich countries is that large revenues may make spending decisions worse from a social point of view (Ades and Di Tella, 1999; Tornell and Lane, 1999). Adverse partisan influence over these decisions is often hard to avoid. 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. There is substantial evidence that government consumption in developing countries is pro-cyclical and that this increases business cycle volatility (Ilzetzki and Vegh, 2008). This is likely to have negative effects on growth, in particular in developing countries with less developed financial markets.[2] A related issue is the selection of specific investment projects. Earlier research has documented that spending of resource revenues quite often goes toward projects with low returns, motivated by pleasing various political groups or electorates (Murphy, 1983; Gelb, 1988; Little et al., 1993; Isham and Kaufmann, 1999). This is related to the time profile of spending as the capacity to absorb large funds for investment is often insufficient in developing countries (van der Ploeg and Venables, 2012), which calls for postponing the domestic usage.

Given these potential negative effects of high spending, for a developing country such as Uganda, it may be important to either constrain the extraction of oil or the government's usage of the proceeds (see, e.g., Segura, 2006; Davis et al., 2003; Barnett and Ossowski, 2003).[3] However, this is of course inconsistent with the recommendation that a poor country should increase its current spending in order to mitigate capital scarcity and low consumption by the current generation, as discussed above. Which of these two contradicting aspects should be prioritized by developing countries remains an open question. To take a step towards answering this normative question, we analyze quantitatively the potential loss of constraining aggregate spending from resource revenues. Spending can be constrained in two ways: either by 1) postponing extraction or 2) by postponing the usage of the proceeds.

The first option, of postponing extraction, will in most realistic circumstances reduce the total economic value of the oil. We analyze how large, quantitatively, the potential losses are of deviating from the profit-maximizing extraction path. We find that postponing extraction will generally entail large losses relative to the maximum value that can be obtained. Hence, this seems to be an unattractive solution to the problem.

The second option, to constrain government spending, can be done by investing the resource revenues in a sovereign wealth fund, and utilizing a simple and transparent rule that stipulates annual spending to be a fixed share of the fund's total value.[4] While such a setup may not always

1. While Uganda, specifically, seems to be rather stable, there are worrying signs where the process around exploration rights lacks transparency and the appropriate checks and balances. An example of lack of transparency is the sharing agreements of oil revenues in Uganda: they are not public. This is likely to limit the ability of the media and the citizens of Uganda to scrutinize the agreements, investigate whether they contain problematic elements and whether they address all the important issues in a proper way. Furthermore, the recent oil bill that was approved by parliament and the president of Uganda (EDP Bill, 2012) gives the Minister of Energy and Mineral Development the authority to sign and revoke agreements with oil companies. To avoid 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, approved by parliament and executed by an independent authority.

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

3. Even though part of the motivation for this recommendation tends to lie in the theoretical explanation forwarded by Friedman's permanent income hypothesis, the concern about negative political consequences is also cited as an important justification for constraining spending.

4. In Uganda, following the oil bill (EDP Bill, 2012), a sovereign wealth fund is supposed to be set up. However, the parliament each year decides how much of the oil revenues will be invested in the fund and how much will be spent right

be in the interest of politicians and elites, it is, under some circumstances, a realistic outcome. A simple and rigid framework, such as a fund and a spending rule, partly ties the hands of future politicians by making the breaking of the rule easy to detect for media, international bodies and non-incumbent politicians (Bacon and Tordo, 2006; Engel and Valdés, 2000; Landon and Smith, 2015). In particular, it is simpler to implement than a theoretically optimal rule that takes into account all future revenue flows, since this requires forecasting the notoriously unpredictable oil price (Landon and Smith, 2015). The uncertainties in such predictions may undermine transparency in spending decisions which is likely to create a bias towards spending and unsustainable borrowing by incumbent politicians. Moreover, to the extent that current politicians can constrain the actions of future politicians (through, say, constitutional mechanisms), it may be in their interest to set up a spending rule to minimize over-spending by their opponents in the future. A spending constraint can be implemented in more or less binding ways. For instance, as a consensual hand-shake between parties in parliament (as in Norway) or in more binding fashion through formal rules. It should be emphasized that it is the constraint for how much has to be put in and how much can be taken out from the fund that mitigates partisan politics and other political effects. It is not the fund as such.

Clearly, for a capital-scarce and borrowing-constrained developing country, where the marginal utility of consumption by the current generation is relatively high and investment needs are many, constraining current spending through spending rules involves welfare losses. But how large are these losses? To answer this question, we build a macroeconomic model which is well suited for a quantitative analysis. Our framework is quite standard in macroeconomic analysis of capital scarcity and borrowing constraints (see Gourinchas and Jeanne, 2006). It contains the intra-temporal trade-off between public and private capital and the intertemporal trade-off between capital and consumption. It further allows for population growth, investment frictions, capital scarcity, borrowing constraints and technical change—features that have significant relevance for developing economies. Thus, our paper departs from other studies that examine welfare implications of spending rules in the context of advanced economies where these features of developing countries are abstracted from (see, e.g., Landon and Smith, 2015).[5]

Perhaps surprisingly, in the baseline calibration of the model (abstracting from any political side effects), we find that from a consumption smoothing perspective, a shift from a spend-as-you-go scheme to using an oil fund along with a fairly strict spending rule appears to entail only a marginal, if any, welfare loss. In particular, the losses appear small compared to losses of fairly mild political side effects that may arise as a result of increased oil spending. For instance, if the spend-as-you-go scheme implies a lag of structural transformation by one year, or if it retards annual productivity growth by as little as 0.006 percentage points, then the fund is preferable. This implies that, considering the potential negative political and economic side effects of a drastic increase in oil spending, the case for constructing a sovereign wealth fund along with a spending rule is rather strong. Furthermore, the most serious political side effects of resources, such as civil conflict as suggested by Collier and Hoeffler (1998) and Garfinkel and Skaperdas (2007), should imply welfare losses of several orders of magnitude larger (see, Abadie and Gardeazabal, 2003). Thus, to the

away. Likewise, the parliament decides how much 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 briefly discuss as well. However, it may also create the sort of problems discussed above. A spending rule can mitigate these problems.

5. Landon and Smith (2015) analyze volatility in the context of a *developed* country. For a detailed review of the literature on volatility in developing countries, see Loayza et al. (2007).

extent that a spending constraint can reduce the risk of these events, it would be well worth the small loss of consumption smoothing.[6] Put differently, getting the political checks and balances right is more important than getting the spending profile right.

Our paper relates to a number of studies that derive optimal spending rules for resource-rich developing countries (see, e.g., Berg et al., 2013 and Cherif and Hasanov, 2013). However, our approach and question is largely different. The earlier studies focus on deriving the optimal spending profile assuming specific resource-curse problems; such as resource income volatility (Cherif and Hasanov, 2013) and Dutch disease (Berg et al., 2013). These studies are rich in details and therefore provide important quantitative insights on the workings of the specific resource-curse mechanisms they model. However, it is difficult to know how the optimal spending rule changes if one would incorporate other mechanisms which are harder to quantify (such as conflicts and corruption). Hence, by abstracting from any one specific resource-curse mechanism, we provide what could be considered as an upper-bound estimate of the welfare losses from constraining spending. This upper bound can then be compared to any resource-curse mechanisms, which are typically difficult to quantify in detail in a unified framework but for which one may have a " ball-park" estimate of the losses. Our results show that the welfare loss from constraining spending is relatively small.

The next section presents an overview of the Ugandan economy and the size of the oil revenue. This is followed by analysis of how quickly to extract the oil (in Section 3). Sections 4–6 use a macroeconomic model to analyze different spending schemes. Section 7 concludes.

2. SOME KEY CHARACTERISTICS OF THE UGANDAN 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.[7] Using a purchasing power parity (PPP) adjustment to take a lower domestic price level into account, income is almost three times as high at a GDP per capita of \$1,345.[8]

Like many other countries in Africa, growth in recent years has been relatively high.[9] Figure 1 presents the growth rate of PPP-adjusted GDP per capita in constant prices for the period 1975–2010. The solid line depicts the average growth rate over the preceding ten years while the dotted line depicts the yearly growth rate. As can be seen, the average growth rate was 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%. Over the last decades, the population growth rate was around 3% annually, implying an average total GDP growth rate of around 6%.

Government finances have been in deficits in recent years, but over the last decades, deficits have not been very large and the general government gross debt stands at 36% of GDP in the year

6. Furthermore, using Sachs and Warner (2001) results and applying them to Uganda's expected oil income, suggests a negative effect on growth of 1 percent annually.

7. GDP levels for 2011 are taken from the World Bank at http://databank.worldbank.org, accessed in September 2013.

8. Uganda has a GDP per capita that is slightly lower than those in neighboring Kenya and Tanzania, but substantially higher than those in South Sudan and the Democratic Republic of Congo.

9. Note, however, that there are uncertainties regarding the data quality. There is, in fact, a suspicion that growth and GDP levels in sub-Saharan Africa are seriously underestimated; see (Young, 2012). Needless to say, our quantitative analysis is conditional on the data being accurate.

Figure 1: Real Per Capita Growth in Uganda—Year-to-Year Growth Rate (dotted line) and the Average Growth Rate Over the Preceding Ten Years (solid line)

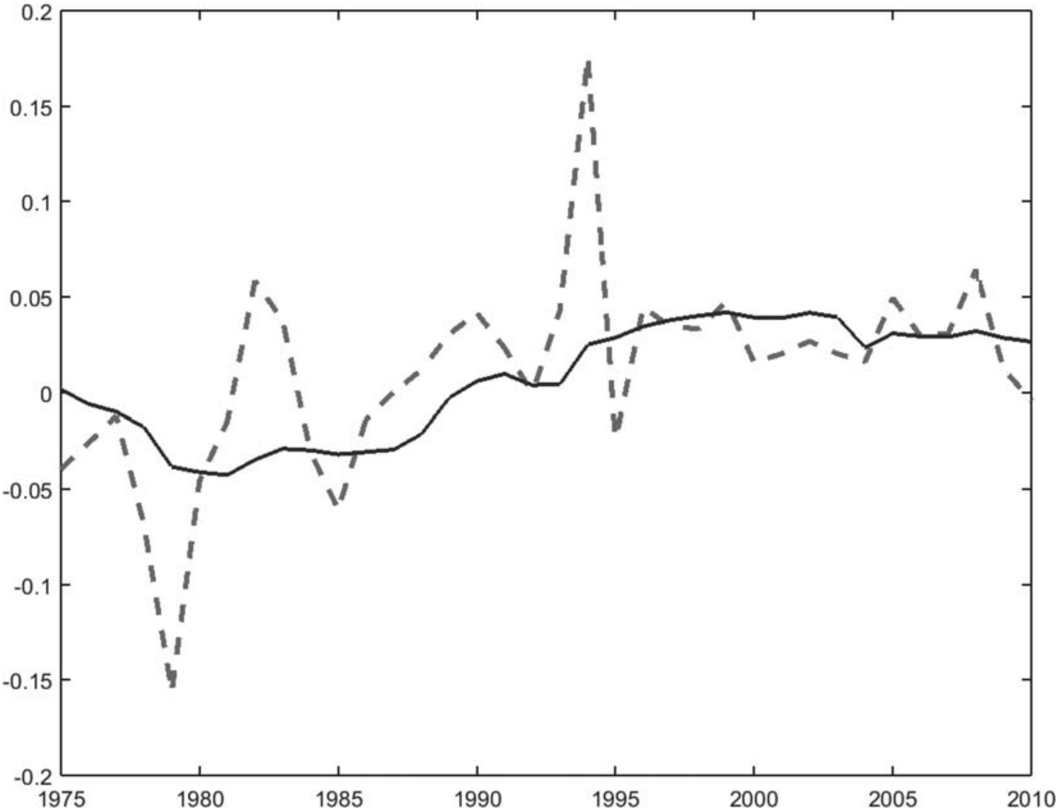

Source: Penn World Tables 7.1.

2012. However, the current account deficits have recently been deteriorating at an alarming rate, as can be seen in Figure 2.

2.1 TFP Decomposition

Let us now perform a simple decomposition of Ugandan per-capita growth into the contribution of capital accumulation and productivity growth. This decomposition will be used when calibrating the model we develop later. The approach is standard using a stylized production function of Cobb-Douglas type, i.e.,

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

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 in period t. The exponents γ and α are time-invariant parameters. Using data on output growth and how the capital stocks and labor evolve over time, we back out how much output growth can be accounted for by growth in factor inputs. The remainder is then

Figure 2: Current Account (solid) and General Government Net Lending (dashed) in Percent of GDP

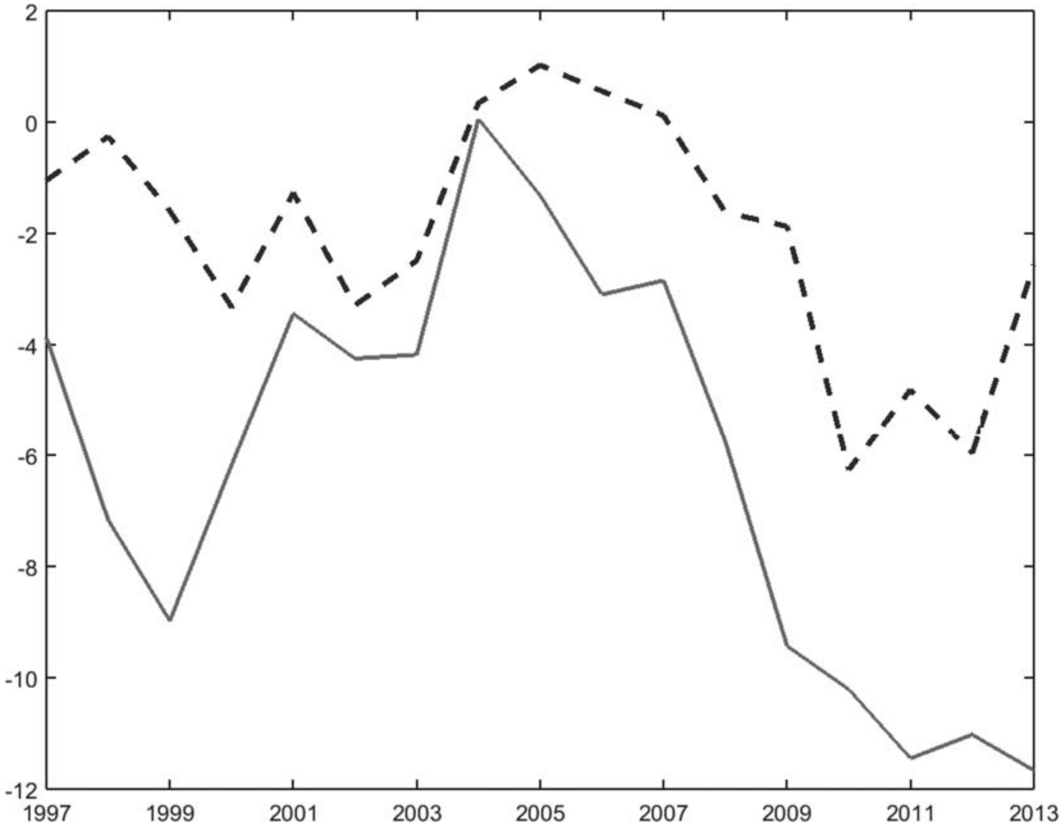

Sources: IMF World Economic Outlook, October 2012 Online edition.

attributed to TFP growth. This procedure would be standard if one abstracted from government capital. The appendix outlines the procedure details.

Figure 3 presents the yearly growth rate of TFP. As in the previous graph, we show the average over the preceding ten years. As can be seen, the TFP growth rate increased substantially between the 1980s and the 1990s—from negative numbers to levels between 2 and 3%. However, there is a clear sign of a productivity slowdown and, over the last 15 years in the sample, the average productivity growth rate was a fairly low 0.65% annually.

We can further conclude that productivity growth contributed to slightly 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 yearly growth in per-capita GDP was 3.9%, of which 2.4 percentage points are accounted for by productivity growth. Over the next decade, GDP per capita grew by 2.7% annually, of which only 0.1 percentage points are accounted for by productivity growth.

The estimation procedure also allows for balanced growth—a situation where investment rates are constant while per-capita GDP and capital grow at the same rate. Such a balanced growth path is a key feature of the model presented in Section 4 since the economy tends to converge to such a growth path whenever the productivity growth rate and other parameters of the economy

Figure 3: Ten Year Backward Average of TFP Growth Rates in Uganda

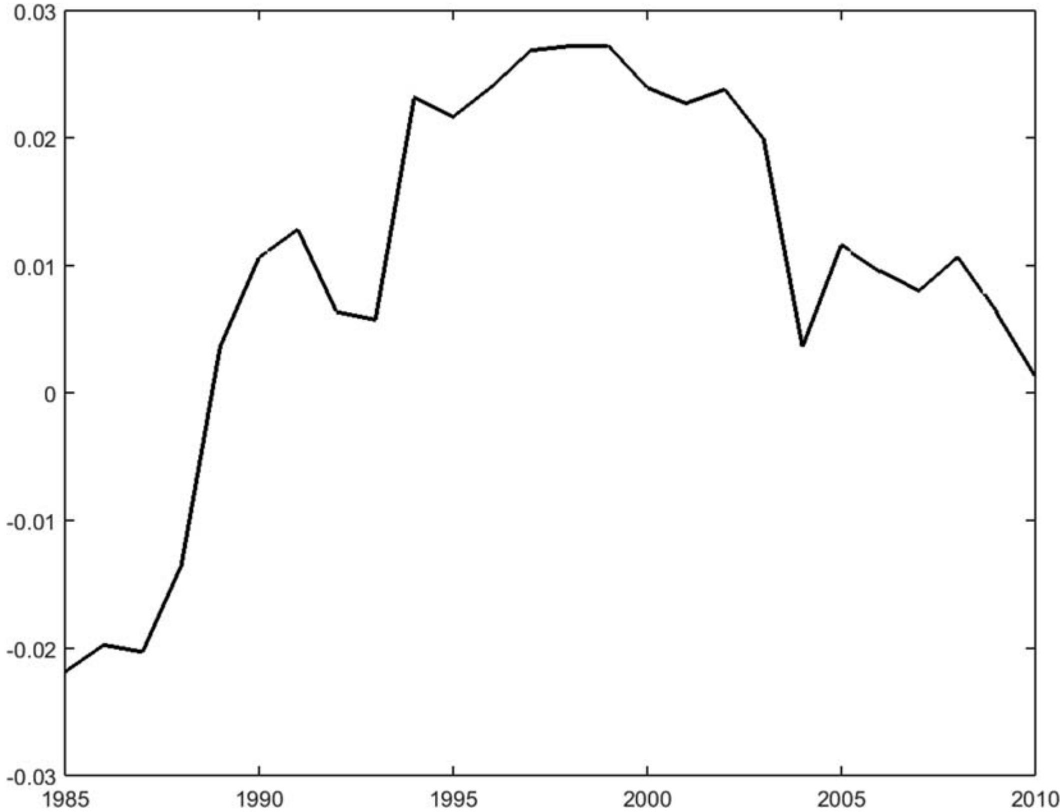

are stable. In conclusion, GDP growth was close to balanced during the 1990s but unless productivity growth picks up and an important structural change occurs, one might worry that growth will fall. It is important to note that the oil discovery, however positive as a source of wealth, is in itself not an automatic source of structural change.

2.2 Oil Resources

There is substantial uncertainty regarding the amounts of recoverable oil in Uganda. Over one billion recoverable barrels have been discovered in the Lake Albert Rift Basin.[10] There is also potential for more discoveries. According to the organization Oil in Uganda, a reasonable estimate is 2.5 billion barrels.[11] Following a discussion with Tullow Oil during our field visit to Uganda, we use as a somewhat conservative benchmark assumption that 1.8 billion barrels 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).[12] This sums to roughly $15–

10. Stated on the web site of Tullow oil, http://www.tullowoil.com, in February 2013.

11. Stated on the website of Oil in Uganda, www.oilinuganda.org, in February 2013.

12. This assumes that an efficient infrastructure for transportation is constructed in Uganda (e.g., a pipeline). The pipeline has been one of the chips in negotiations between the Ugandan government and the oil companies (we return to this later). Transportation by truck is likely to be substantially more costly at around $23 per barrel (Henstridge and Page (2012).

$20/barrel of costs. This is not an exact calculation of costs but suffices for producing an estimate of the economic magnitude of the oil resources.

With a Brent price around 65 US$/barrel, these figures would mean profits of around $45/barrel.[13] At those prices and costs, the oil resources amount to $99 billion in revenues and $81 billion in profits. With a population of 35 million, the latter means $2300 per capita. Note that this assumes all oil profits end up with the Ugandan government while the oil companies get none of the resource rents. This is obviously unrealistic as the sharing agreement probably specifies some rents also to the oil companies. However, as the sharing agreements are not public and the government officials did not share them with us (we see the confidentiality as a problem in itself), we make this stark assumption here and treat the value of $2300 as an upper bound. This is clearly a sizable amount as compared to a per-capita income of $487 per year.[14] However, since GDP is a flow, it may be more reasonable to compare it to the income flow from oil. As a thought experiment, suppose that Uganda invested all oil profits in a sovereign wealth fund which yields a real rate of return of 4% per year. This would mean an added income flow per capita of $92 per year. That is, an increase amounting to 19% of current income could be sustained forever. Another perspective of the value of the oil resource is obtained by noting that with a growth rate of 3% per year, a 19% higher level of income is achieved after about six years. The purpose of this back-of-the-envelope calculation is to show that the value of the discovered oil is not large enough by itself to lead to a revolutionary change in the living conditions of the average Ugandan. However, wisely spent, the oil income can certainly make a difference.

3. WHEN TO EXTRACT

In the previous calculation of the total value of the oil resource, it was assumed that the oil could be extracted at once. Although it is useful as a first approximation, this assumption is of course unrealistic. 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 technically feasible extraction profile. The solid line in Figure 4 shows this profile measuring extraction in 1,000 barrels per day. We take it that Tullow has incentives to propose a high extraction speed; while there may be social reasons to slow down the extraction, it is hard to see such motives for a private international oil company. Thus, we take the profile in Figure 4 to represent the quickest possible extraction—an upper bound of sorts. As the figure shows, 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.

Before extraction starts, it is to some extent possible to alter the speed of extraction by choosing different capacities of extraction and transportation equipment. Once the investments have been made, however, such flexibility is very costly.[15] We will now calculate the present discounted value of oil at two different extraction profiles—the quick extraction path of Tullow and a path that

13. The quality of the oil is somewhat low; hence the actual price should probably be slightly below the Brent crude oil index. We abstract from this in the analysis. We use 65 $/barrel as our base case since this roughly corresponds to the sustainable extraction cost of shale oil which is believed to largely determine the price of oil today and over the next decades.

14. 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.

15. The speed of extraction has been one of the main points of disagreement between the oil companies and the Ugandan government where the government advocates slower extraction.

Figure 4: Tullow Extraction Profile in 1000 Barrels Per Day (solid line) and Constant Extraction (dashed line)

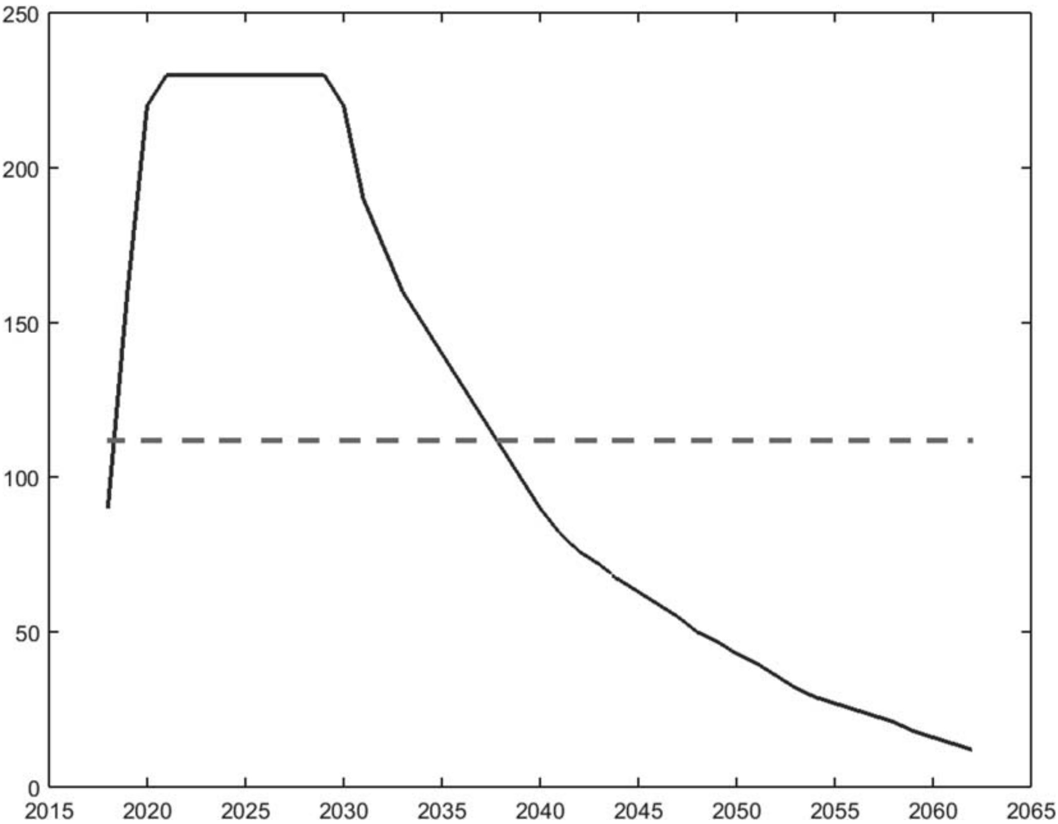

is constant over the same 45-year period—and at two different oil price scenarios—a gradually increasing price and a constant price. By extracting 112,000 barrels per day, the same amount is produced over a 45-year period as under the Tullow plan (that is, 1.84 billion barrels). This extraction profile is represented by the dashed line in Figure 4.

The present discounted value of the oil is given by the expression

$$W = \sum_{t=0}^{T} \frac{1}{R^t}(p_t E_t - C(E_t)).$$

It simply says the net profits (revenues net of extraction costs) in each period t should be discounted by the accumulated interest from today until that date (R^t). Hence, to calculate the value of oil, we need to make assumptions about future oil prices, the accumulated interest, and the evolution of costs. We take as a benchmark that the (real) interest rate is 4% per year—a standard assumption in macroeconomics. We further assume that the costs are constant and examine different assumptions about the price path. As a first alternative, consider the common assumption in theoretical resource models, that the Hotelling rule (see Hotelling, 1931; Dasgupta and Heal, 1974) applies. This theoretical result states that if resource owners of a finite resource are rational, the path of equilibrium aggregate supply should lead to growth of a resource price (minus extraction costs) at the same rate as alternative investment opportunities. The result stems from owners treating the resource as an asset. By not selling it, the returns they get is in the form of price growth which has to equal the interest in order for returns on assets to equalize. Thus, one scenario is that the oil price grows at the rate of the alternative investment opportunity, in our case 4% annually.[16]

Under the assumption that oil price growth is equal to the discount rate, starting from an oil price of $65 and that extraction follows the Tullow profile, the present discounted value of oil revenues minus costs is $84 billion. Instead, assuming a flat extraction profile, the value is $93 billion, i.e., almost the same. In fact, if we consider the assumptions underlying the Hotelling result, the timing of extraction has no effect on the value of the oil resource since the price then rises at the rate of interest.[17] Under alternative, and possibly more realistic, assumptions about the future oil price 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 an alternative oil price scenario—one where 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 at $65, the value of the Tullow extraction profile is $52 billion, while the flat extraction profile gives a value of $40 billion. Thus, the delayed extraction profile implies a loss of $12 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 and hence undesirable. In fact, the cost of delaying the whole extraction profile by one year equals the discount rate times the total value when prices are assumed to be constant. That is, 4% of $40 billion which equals $1.6 billion. It is worth mentioning here that a delayed extraction profile also makes the value of the oil resource more sensitive to variations in future oil prices.

These losses of delaying extraction can be used to put one of the negotiation problems in Uganda into perspective. In Uganda, the government may delay building the pipeline (and hence the start of extraction) since it wants an oil refinery to be built first (Petroleum Intelligence Weekly, 2015). The refinery is meant to supply the Ugandan and neighboring markets with products derived from oil. While this may or may not be a good idea looked upon in isolation, it is important to verify that the potential benefits of such a refinery are greater than the cost of delaying extraction which, as was just shown, can be substantial.

4. A MACROECONOMIC MODEL OF INTERTEMPORAL TRADE-OFFS

This section provides a formal analysis, using a calibrated theoretical model, of the role of oil revenues in the Ugandan economy. We take a growth perspective, since the main issue is the role of oil revenues over time. Thus, the maintained time horizon is rather long—in the 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 catch-up via a gradual, but slow, elimination of inefficiencies. This approach is consistent with the recent growth literature as, for example, exposited and summarized by Jones (2013). The removal of inefficiencies implies direct benefits to output but also

16. While the oil price has been falling during most of the 20th century and price increases (such as during the oil crisis in the 1970's) have been followed by subsequent price falls, one possible scenario is that the increasing price trend observed over much of the period after 1998 will continue in the long run. See, for instance, Spiro (2014), Venables (2015) and Hamilton (2009), who analyze the long-run oil price and Hart and Spiro (2011) for a discussion.

17. The 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.

indirect benefits through induced capital accumulation. The model displays this process and how oil revenues arriving in different years affect it. The model also incorporates capital scarcity and limited access to borrowing—important features of developing economies. This is consistent with the recent literature on the macroeconomics of developing countries (Gourinchas and Jeanne, 2006).

We abstract from a number of issues. For instance, business-cycle and other fluctuations and the often-discussed Dutch-disease problems arising from resource-income shocks are not studied in the formal analysis. Furthermore, we use a representative-agent framework which implicitly abstracts from distributional issues (of oil and other income). The concluding discussion in Section 7 remarks on their potential impact.

We will begin by describing the basic features of the model and the main quantitative assumptions (on parameter values and the oil revenue). Then, in Section 5, we present a sequence of illustrations of the dynamic paths of major variables from the calibrated model and show the effect of oil revenues on these paths. Comparison of welfare under alternative spending schemes will be presented in Section 6.

The baseline setting of the model abstracts from political-economy factors. But in evaluating welfare we add three possible political side effects of high oil spending and examine how they impact the welfare comparisons. We also report on a number of robustness checks to the baseline setting of the model.

We use an extension of the basic Cass-Koopmans model of optimal growth, the cornerstone of growth theory (Koopmans, 1963; Cass, 1965). This model focuses on a macroeconomic aggregate output which can be used for investment or consumption. The main model describes Uganda as a closed economy with no access to borrowing. While, in reality, Uganda is probably not completely credit constrained, it seems such constraints are strong enough to motivate no access to borrowing as the main case. However, we will also present the paths of major variables under the hypothetical case of Uganda being a completely open economy. This case is not a realistic scenario. However, it serves as a way of gaining intuition on how partial relaxing of the credit constraint may affect the economy. We also briefly discuss how different degrees of openness would impact the comparison between the two alternative schemes of using the oil revenues. It should be noted that if Uganda had perfect access to credit markets, the economy's path would be independent of the choice of spending scheme. This happens because the overall spending in the economy will ensure that returns in Uganda would always be equal to the returns in the rest of the world. It therefore follows that any extra investment coming from oil income would fully crowd out other investments, making the spending profile of oil revenues irrelevant for the dynamics of the economy. Thus, the fact that Uganda's economy, and most developing countries with it, are not perfectly open is what makes the exercise of comparing different spending schemes interesting.

When oil is included in the analysis, extracted oil is added to the aggregate income. In all versions of the model it is assumed that oil is traded internationally at competitive prices.

The budget constraint is

$$c_t + i_t + i_{gt} = y_t + p_t o_t, \tag{2}$$

where c is aggregate 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 put 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 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 that of Uganda; the infrastructure necessary for private production is only partially in place, 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

$$k_{t+1} = (1-\delta)k_t + (1-\kappa_t)i_t \tag{3}$$

and

$$k_{g,t+1} = (1-\delta_g)k_{gt} + (1-\kappa_{gt})i_{gt}, \tag{4}$$

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. Thus, we take the perspective here that when one unit is invested, a fraction disappears. This is motivated by discussions with Ugandan officials and by recent research (van der Ploeg and Venables, 2012). 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

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

We thus assume a Cobb-Douglas production function, giving constant shares to capital and labor income.[18] The labor input is suppressed (i.e., it can be viewed as set to 1 and to be inelastic). Population growth is not modeled in this baseline setup—an issue whose implication will be highlighted later on. Government capital, as modeled, displays complementarity with private capital: with a higher stock of government capital, 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 are thought of as 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 endogenous capital accumulation. Thus, as with the lowering of the κ's, improved efficiency has both direct and indirect benefits. Note, finally, that we abstract from oil as an input in production, since oil's primary importance is as a source of extra revenue. This is in line with the approach taken in the literature on economic growth.

We assume a representative agent with preferences

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

18. We change notation slightly relative to Section 2 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 (6). This change is for convenience only and does not affect the analysis.

where $u(c)=\dfrac{c^{1-\sigma}-1}{1-\sigma}$. 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 the individual wants to smooth consumption over time. T is finite but large enough for the economy to have converged to a balanced-growth path. To find an optimal path, we thus maximize the above utility function under the restrictions implied by the macroeconomic resource constraint (2) and the capital accumulation equations (3) and (4).[19]

In the baseline calibration of the model (see the appendix for details), we view the past as quite distorted both in the private and public investment sectors, though more in the latter. Based on discussions with Ugandan officials, we set them to be initially 25% and 50% respectively.[20] However, these distortions, as described by the κ's, are assumed to asymptotically vanish. To account for a catch-up effect, we assume total-factor-productivity growth to be 4% initially (to capture the growth experience in the last twenty years) and then declining toward 2.5% (a rate consistent with the rest of the world). Asymptotically, the closed economy will generate the same interest rate as in the rest of the world but in the short run, due to borrowing constraints and capital scarcity, the interest is higher in Uganda than in the rest of the world. For the parts of our baseline economy not having to do with oil, we mostly use standard parameter values in the macroeconomic literature (see the appendix for details).

As for oil revenues, we consider the baseline scenario as the "maximum extraction path" suggested by Tullow oil (see Figure 4). This path implies 33 million barrels extracted in an initial year of 2018 and then a rise to a maximum of 84 million in just a few years (obtained in 2021). The high level of production is maintained for nine years and then it gradually declines with the last oil production occurring in 2062. We use \$65 per barrel as the baseline price and assume this price (in real terms) will be constant.[21] 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 (see Section 2.2). These costs are assumed to be constant over time in real terms.[22]

A final 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 oil price of 65, this size is obtained by selecting the barrel unit relative to total output appropriately. Relative to GDP in 2011, this gives 8.8%.[23] It should be noted that we use official GDP figures and a current exchange rate, as opposed

19. 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.

20. These values are also roughly consistent with Pritchett (2000) and Hurlin and Arestoff (2010) who estimate investment inefficiency in various Latin American and African countries.

21. A number of alternative price scenarios have been analyzed for robustness. We use a constant price of \$65 per barrel since this roughly constitutes the extraction costs of shale oil, the supply of which can be expected to last for at least a decade or two.

22. As discussed in Section 2, the sharing agreements are not public. Therefore, we assume that all the oil profits accrue to the government and abstract from the sharing of profits with the oil companies. This is, however, not likely to significantly affect the main macroeconomic dynamics as sharing the oil revenues would simply imply a proportionate income effect in all periods. The same of course holds if we lower the oil price proportionally in all periods. We discuss the robustness of the results in this respect when comparing welfare in the next section.

23. Uganda's GDP in 2011 was measured at \$16.8 billion. Tullow's estimate for 2018 of 32,850,000 produced barrels means a revenue of $(65-20)\cdot 32.85\cdot 10^6$, which equals \$1.48 billion.

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.[24]

5. MACROECONOMIC DYNAMICS

We will now present a sequence of illustrations of the dynamic paths of major variables from the calibrated model, beginning with the case where Uganda proceeds without any oil income. This case is not interesting per se but it serves as a comparison and as a way of explaining the basic features of the economy.[25] This case without oil is illustrated for both a closed (i.e., borrowing constrained) and an open economy. Then, we look at how oil revenues affect the closed economy outcomes under two different scenarios of revenue usage—the spend-as-you-go (SAYG) scheme and the fund scheme. In the SAYG, the oil revenues immediately enter the domestic economy and are optimally divided between consumption and investment. In the fund scheme, the oil revenue is put in a fund from which, importantly, only a limited amount is withdrawn each year and is optimally divided between consumption and investment. Although our analysis could be amended to allow any intermediate use of the oil resources, we believe that these different cases bracket most of the remaining possibilities because they emphasize the possible advantages and disadvantages of different setups.[26] These two cases of oil-revenue spending are then compared to an open economy with unlimited access to borrowing and lending.[27] This comparison with the open economy illustrates the implications of credit market constraints for the macroeconomic dynamics under the SAYG and fund schemes; thus, the comparison highlights how the credit constraints may affect the costs and benefits of the two schemes. In all cases, we simulate the model for a hundred years starting from 2013.

5.1 The Case without Oil

Under the assumption that Uganda receives no oil revenue, the economy should be expected to converge rather smoothly to the balanced growth path, as TFP growth gradually slows

24. The model here does not have the distinction between traded and non-traded goods.

25. The solution technique relies on global, nonlinear solution methods to solve a series of second-order difference equations (in k_t,k_{t+1} and k_{t+2}). We apply the shooting algorithm where, given k_0, we search for k_1 so that the sequence $\{k_t\}$—a set of solutions from the second-order difference equations given k_0 and k_1—converges to the steady state value of k. The programs are written by the authors and are available upon request.

26. We also looked at a case where a constant amount of oil money is spent every year. This case would necessitate borrowing in early periods against the future oil revenues, i.e., this rule takes into account the oil under ground. In terms of welfare, this rule is quite similar to the spend-as-you-go scheme and we therefore leave it out from the figures and detailed analysis. As argued in the introduction, this rule relies on forecasts of the oil price over the next 50 years and the quantity of oil in the ground. The significant uncertainty surrounding these quantities is likely to open up a substantial room for political discretion (in estimating/forecasting quantities), and can be used by incumbent politicians to maneuver spending. Hence, the constant spending rule may run the same kind of political risk that one wants to avoid in the spend-as-you-go scheme.

27. The economy that is open to international borrowing and lending has a resource constraint that reads $c_t + a_{t+1} + i_t + i_{gt} = y_t + p_t o_t + (1+r)a_t$, where a denotes transformed international lending 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 $a_{t_{t=0}^{\infty}}$, and the “ no-Ponzi” condition.

Figure 5: Capital, Output and Consumption in a Closed Economy, No-Oil Case

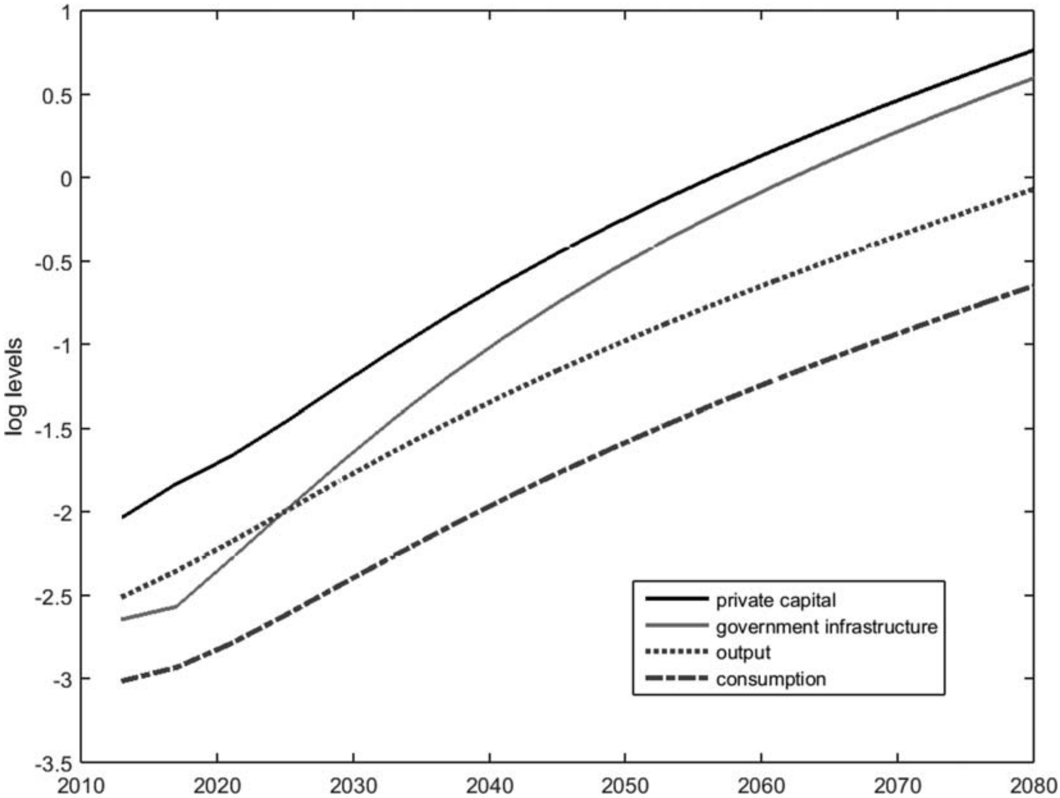

down and as the investment wedges disappear. Figure 5 displays the paths for (logarithms of) the major variables.

We see that all variables grow roughly 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 assumed to be larger than for private-sector capital. We see that the transition path is somewhat non-smooth 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.[28] 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 revenues that we look at below.

An important feature of developing economies such as Uganda is scarcity of capital. As can be seen in Figure 6, this feature of capital scarcity is reasonably captured in our model. The steady-state private capital-output ratio is about 2.4, to which the economy converges in the long-run. However, during the transition, the calibrated capital stock lies significantly below the steady-state level for an extended period of time. For example, during the first ten years, the calibrated private capital-output ratio is about 1.7, which is only 70 percent of its steady state level. Note also

28. 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.

Figure 6: Capital-Output Ratios and Marginal Products of Capital in the Case without Oil

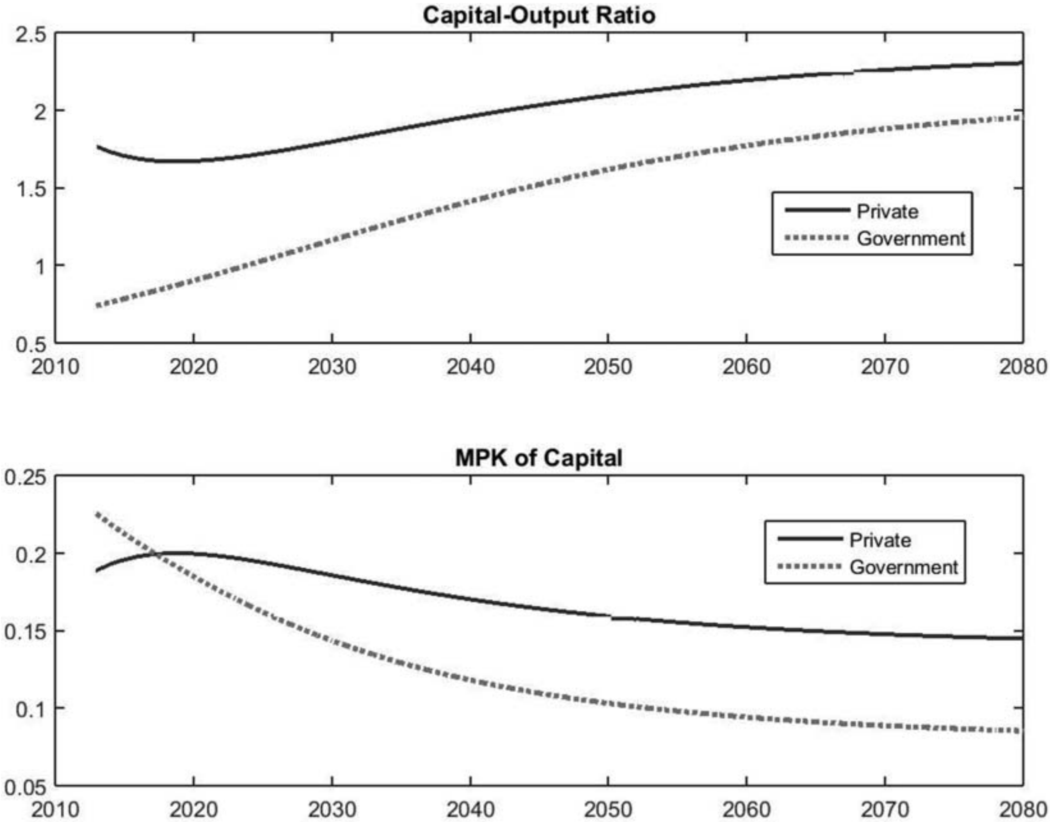

that this number is quite close to Uganda's actual capital-output ratio over the past decade, which averaged about 1.6. The government capital-out ratio is also below the steady-state level for an extended period of time.

The low capital-output ratio during Uganda's transition to the steady state means that the marginal product of capital is quite high (see bottom panel of Figure 6). This is true both for private and government capital. The marginal product of government capital during the earlier periods is more than twice its long-run value. During its peak, the marginal product of private capital exceeds its steady-state value by about 50%.

The above paths reflect optimal transition dynamics for a closed economy. If the economy would be open and Uganda could borrow and lend freely at an international interest rate, the dynamics would 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. That is, it is informative for understanding how capital scarcity and borrowing constraints affect a developing country. Figure 7 displays the results for output and consumption, in comparison with the closed-economy case. Several points are noteworthy here. First, while the balanced growth path for output will be the same,[29] the open-economy output would be

29. The model is calibrated so that the real interest rate in balanced growth will be the same.

Figure 7: Output and Consumption in the Closed and Open Economies, No-Oil Case

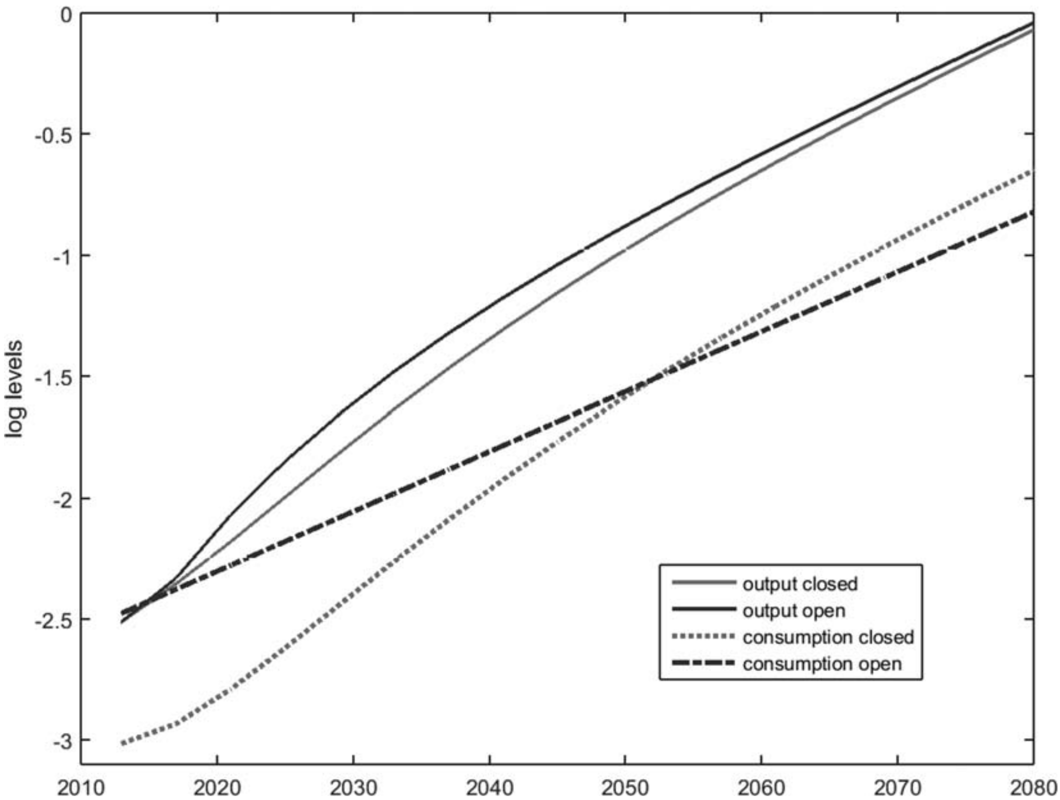

above closed-economy output throughout but with a decreasing gap. The reason is that the internal interest rate in a closed-economy is higher during the transition and only gradually falls to 4%, so less capital is accumulated when foreign borrowing cannot be used to (partly) finance investment. Second, consumption would be quite a bit higher in an open-economy for several decades. This is because Ugandan output is low at the outset compared to its future potential, 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 an open economy would therefore be less steep than in a closed economy. But eventually the paths cross, as loans need to be paid back. Third, welfare would be higher in an open-economy case both due to consumption smoothing and higher output.

5.2 The Case with Oil

We now introduce oil revenues into the economy and look at the dynamics of our key macroeconomic variables under three alternative scenarios: a closed economy with a spend-as-you-go (SAYG) scheme, a closed economy with a sovereign wealth fund and an open economy (with unlimited access to borrowing and saving). The purpose is to examine how the oil revenues affect the macroeconomic dynamics under alternative scenarios.

The maintained assumption in the SAYG scheme 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

Figure 8: Consumption and Output in a Closed Economy: Spend-As-You-Go Versus No Oil Case

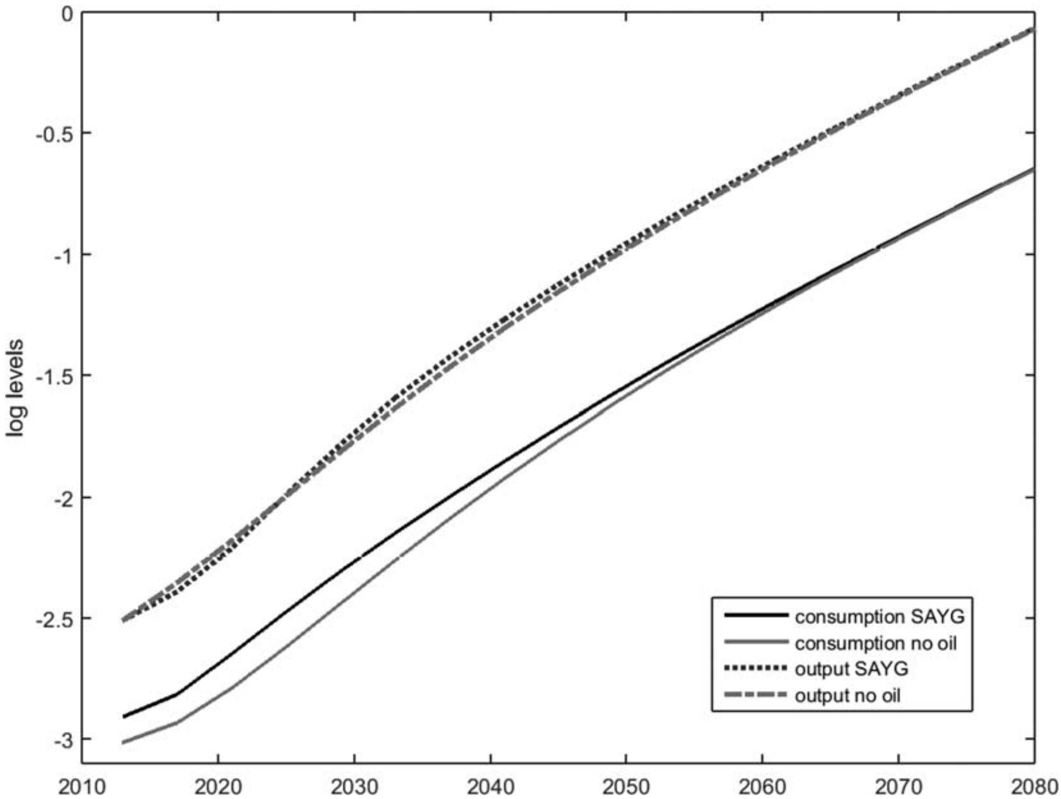

resources are allocated optimally between consumption and investment. The paths for log consumption and output under this scheme are plotted in Figure 8.

We see that, anticipating the oil revenues to arrive later, consumption should be increased already at the beginning. However, facing borrowing constraints, the only way of achieving this is by reducing investments which in turn reduces output in the early years. But then output catches up around 10 years after the first oil revenues have arrived. Output is then permanently higher, though asymptotically the output paths with and without oil are identical.

We now look at the case where a sovereign wealth fund is used. The following fund construction is adopted: (i) extraction according to the Tullow scheme is maintained and the profits are added to the fund as they accrue; and (ii) each year, 4% of the fund is withdrawn for consumption and domestic spending (the balance between which is determined optimally). Figure 9 displays the results for the fund scheme along with the SAYG scheme. As can be seen, the fund construction allows significantly less consumption smoothing than does the SAYG scheme at the same time as it keeps output up in the early years.

If we consider an open economy with oil, the initial consumption response is quite strong: the added oil revenue allows consumption to react even more, since it is possible to borrow against all future oil revenues.

Figure 10 summarizes the results. We compare the consumption and output paths in our three economies with oil. The figure displays the variables (i) in an open economy, (ii) using the

Figure 9: Consumption and Output with Oil: A Fund Versus Spend-As-You-Go

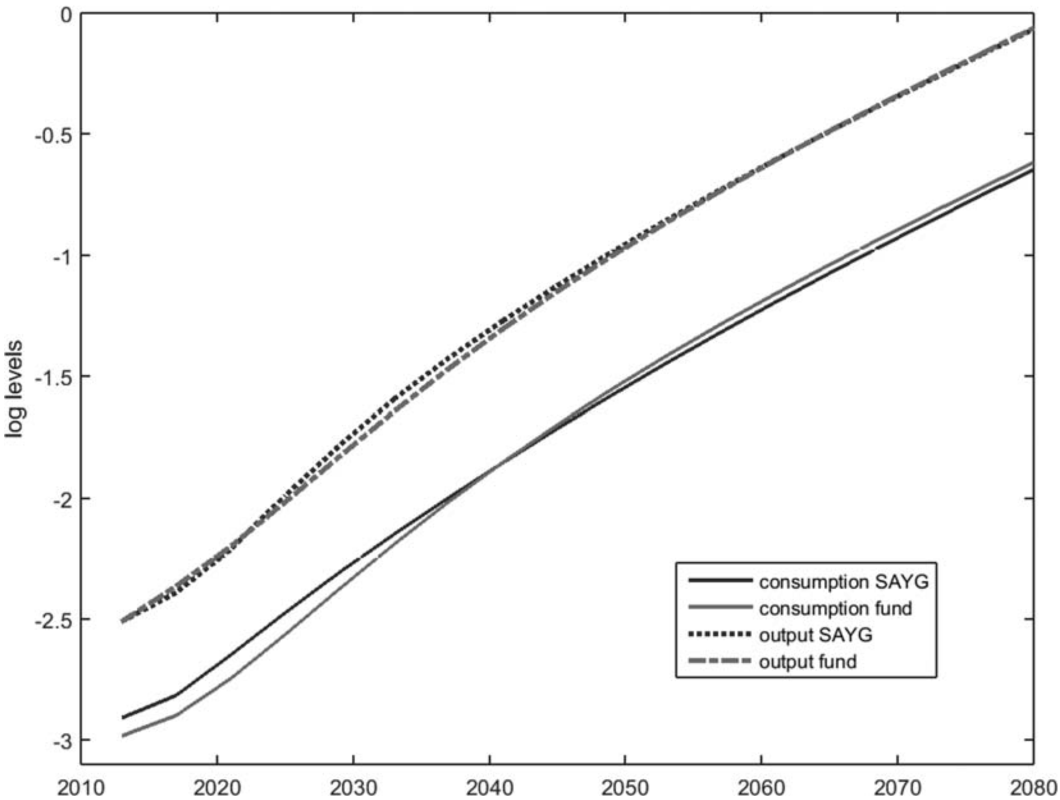

SAYG scheme, and (iii) using the fund and spending-rule scheme. All variables are *relative* to the closed-economy without oil. The upper and lower panels present relative consumption and relative output, respectively.

Under the assumption that Uganda is an open economy, the present value of the oil revenue is annuitized over the infinite future. As mentioned earlier, in a fully open economy, the paths for the macroeconomic variables would be independent of the spending schemes since any domestic spending from oil income would be crowded out by savings in the international capital market. Under both the SAYG and the fund schemes, output is relatively lower. This highlights the possible role of oil revenues in boosting output by mitigating the negative effect of credit constraints on domestic capital accumulation. The rapid increase in output after the year 2020 in the SAYG coincides with the rapid inflow of oil revenues as the extraction pace accelerates, and a significant part of the revenue is spent on domestic capital which, in turn, raises domestic (non-oil) production. Note also that, despite the large amount of domestic spending under the SAYG, output in the SAYG never exceeds output in the open economy. Thus, the domestic rate of return on capital under SAYG still remains below the international rate.

Note that output is lower for many years under the fund. This happens because the fund scheme constrains domestic spending, and capital accumulation is slower than under SAYG after the oil revenues start arriving. Initial consumption is lowest and long-run consumption is the highest under the fund because of the fund constraining initial spending (lowering current consumption) and enabling accumulation of wealth to boost future consumption.

Figure 10: Consumption and Output Levels (Relative to the Closed Economy Without Oil)

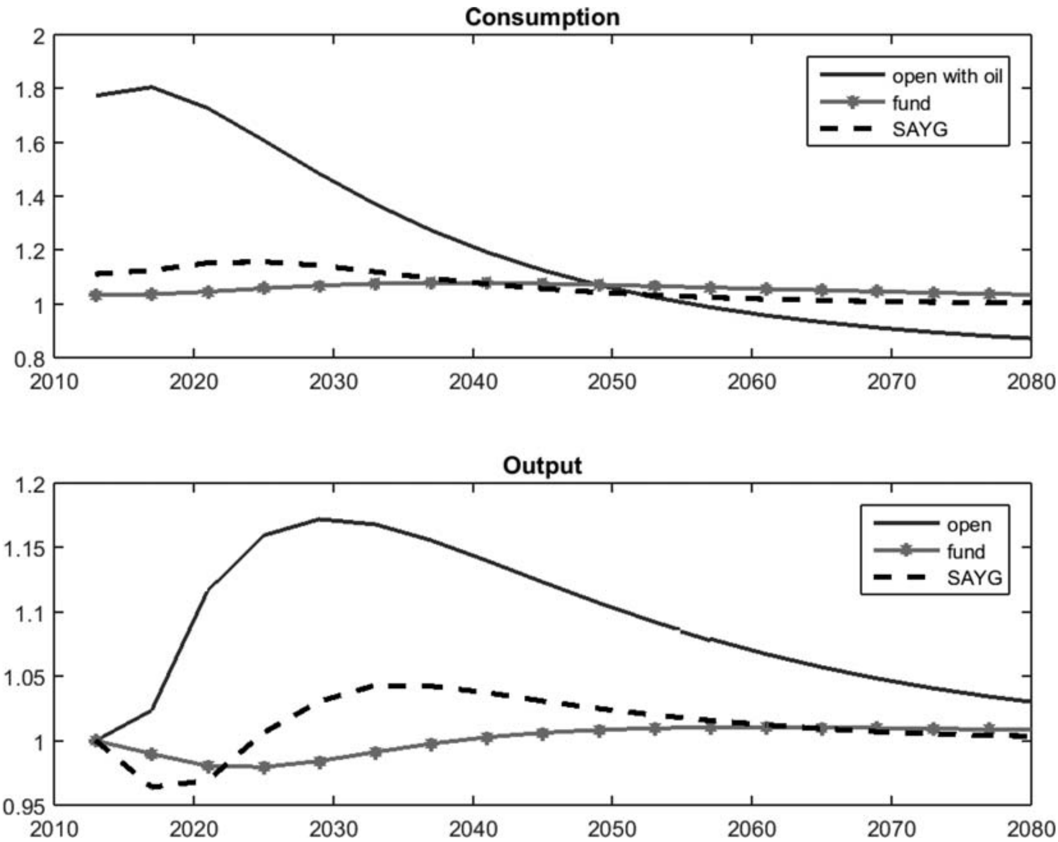

One can think of different degrees of access to credit markets if Uganda is able to borrow at an interest which is higher than the world interest. If, for instance, the interest Uganda faces is higher than the return on domestic investments, then it can be considered as an effectively closed economy. If the interest is lower than this, the degree of openness increases. The comparison between the open economy versus the two spending schemes highlights how access to capital markets may affect the differences in outcomes between the two spending schemes. A relatively better access to credit markets (even though not a perfect one) should make the domestic economy less reliant on the oil spending to finance domestic investment, and the dynamics under both schemes should look more like the ones under the open economy. Thus, in an economy with better access to credit, the difference between the fund and the SAYG is likely to be less significant since the role of oil in accelerating domestic capital accumulation and consumption smoothing would not be as important.

6. WELFARE IMPLICATIONS: A FUND OR SPEND-AS-YOU-GO?

6.1 Benchmark Result

Given that our maintained hypothesis is continued growth and catch-up, with a successive elimination of inefficiencies in production and investment (particularly for government capital), there is a strong need for consumption smoothing. That is, the marginal utility of consumption is initially very high and so is the return to domestic investments and, given the future output increase, restrictions on spending will lower welfare.The implications for oil revenue management are clear, at least on a theoretical level: the oil revenues will help smooth consumption and, foreseeing their arrival, consumption will, and should, rise already today. In a closed economy this is best achieved by the SAYG scheme. In fact, since the domestic rate of return is higher than the international rate, the SAYG is an optimal scheme under borrowing constraints. However, this means that domestic production in the long run is hurt since it lowers investment in the long run.

Considering the case of a fund it is clear that the initial boost in consumption is less marked, as is the initial drop in output growth. The welfare under this scheme will thus, in most settings of our model, be lower than under the SAYG scheme.

How large are these effects, i.e., how large is the loss when keeping initial consumption relatively low under the fund scheme? Arguably not large. For example, in our baseline calibration, the welfare loss from a shift to the fund scheme from the SAYG scheme, measured in terms of a reduction in the permanent consumption level, does not exceed 0.59%.

6.2 Robustness

The result that the welfare loss is low when using the fund scheme (instead of the SAYG scheme) is quite robust to several alternative assumptions. We have assumed a fairly optimistic scenario about the path of Uganda's future productivity and catch-up. There is, of course, substantial uncertainty about this. Thus, as a robustness check, we run the simulations assuming slower TFP growth rates and find that the welfare loss from constraining current consumption under the fund scheme would be even smaller (than in the benchmark case). Our calibration also abstracts from population growth—a relevant factor for Uganda given that it has one of the fastest growing populations in the world. With an annual average growth rate of 3.3 percent, the population of Uganda has nearly doubled during the last two decades—from 17.5 million in the year 1990 to 33.4 million in 2010.[30] Uganda's population growth is also expected to continue at a rapid pace in the coming decades. According to forecasts provided by the UN, Uganda's population may reach 205 million by the end of the century—a sixfold increase as compared to the level in 2010.[31] A rapid population growth can have significant implications for the optimal utilization of the oil revenues. First, the revenues will have to be shared among more people in the future. Moreover, in an economy without international borrowing and lending, higher population growth, all else equal, implies a larger future labor force and hence, less capital per worker for the future generations. Both these factors call for a reduction in the consumption of the current generation. In fact, when we re-calibrate our model incorporating population growth, we find that using the fund scheme delivers a higher welfare than the SAYG scheme.[32]

30. Source: Penn World Tables 7.1.

31. The UN provides three variants of projections for population size: a low, a medium and a high variant. For Uganda, the low, medium and high variants of projections for the year 2100 are, 139, 205 and 292 million, respectively. Source: Population Division of the Department of Economic and Social Affairs of the United Nations Secretariat, World Population Prospects: The 2012 Revision, http://esa.un.org/unpd/wpp/index.htm. Accessed on July 29, 2014.

32. Population growth is incorporated using the standard approach in the Ramsey-Cass-Koopmans model with exogenous population. We take the intermediate population growth projections from the UN (see previous footnote). With c, y, k_g, i_g and i now denoting per capita values, we make the following three sets of modifications to the model. 1) The welfare function is defined as $\sum \beta^t L_t u(c_t)$ with L_t denoting total population size: the welfare of each member within a generation has an equal weight while the welfare of the next-period generation is discounted by β. 2) The resource constraint becomes: $c_t +$

We have presented the analysis for a fund scheme where 4% of the fund is spent annually. However, the relatively small welfare loss from shifting to the fund scheme does not rely on choosing 4% particularly. In fact, given that 4% may not necessarily be the spending share that maximizes welfare, one could minimize the welfare loss further by choosing an alternative spending share with a higher welfare.

Abstracting from the investment wedges (i.e., the κ's) does not change our conclusion. The inefficiencies in investment, as we have modeled them in the baseline case, have two implications. On one hand, a higher level of inefficiency during the initial periods calls for postponing investment to future periods where this inefficiency is lower. This makes the inefficiencies favor the constraining of current spending and adopting the fund scheme. On the other hand, the fact that the inefficiencies gradually disappear in the baseline case makes the future generations relatively richer and induces a motive for increasing initial consumption. This latter effect favors the SAYG scheme. The latter effect of the inefficiencies is, however, quantitatively stronger than the former so taking away the wedges makes the fund more favorable.

In our simulations above, we set the cost of production and transportation to be \$20 per barrel. This is assuming an efficient transport of oil through pipelines, where the transport cost would be \$4-\$5 per barrel. Currently, there are negotiations among neighboring countries to coordinate construction of the pipelines. However, there exists a substantial uncertainty regarding whether the negotiations succeed in putting the pipelines in place.[33] In the absence of an efficient transport system, the cost of transport would increase and hence lower value of the oil. Thus, as a robustness check, we have run the simulations assuming the cost of production to be \$30 per barrel (instead of \$20 per barrel) to account for the possibility of higher transport costs. The welfare loss from shifting to the fund becomes even smaller when allowing for higher transport costs. This is perhaps not surprising since the higher transport cost lowers the value of oil relative to the overall economy, which in turn makes the changes in spending rules relatively less important. This holds more generally for any lowering of the oil value, e.g., if the oil companies get a share of the rents.

Finally, we have also performed the simulation with alternative price scenarios (e.g., a price growing according to the Hotelling model and a constant price at 100 \$/bbl); alternative assumptions for the evolution of the inefficiency wedges κ (e.g., so they do not vanish even asymptotically); alternative values, instead of our benchmark of $\sigma = 1$, for intertemporal substitution elasticity (e.g., $\sigma = 1/2$ and $\sigma = 2$); and alternative values for the discounting β.[34] None of these alternative specifications change the results—the fund yields a loss of at most 2% and usually well below 1%.[35] All in all, this suggests that the potential losses of using a fund are rather small.

6.3 Adding Political Constraints

To put the welfare loss from the fund into perspective, we have analyzed what kind of political side effects would nullify the losses. We will here quantify three alternative externalities of extensive oil spending: i) that it leads to increased waste in investment; ii) that it leads to delayed

$i_t + i_{gt} = y_t + p_t o_t / L_t$. 3) Capital stocks evolve according to $k_{t+1} n_{t+1} = (1-\delta) k_t + (1-\kappa_t) i_t$ and $k_{g,t+1} n_{t+1} = (1-\delta_g) k_{gt} + (1-\kappa_{gt}) i_{gt}$, where n_{t+1} is population growth in period $t+1$.

33. See, for example, Petroleum Intelligence Weekly (2015)

34. Note that this also affects the interest rate, see the appendix for details.

35. The case of $\sigma = 2$ is the one that yields the highest losses, in particular if coupled with a small β (implying more impatience and higher returns on investment also in the long run). There are arguments to be made both in favor of $\sigma < 1$ and in favor of $\sigma > 1$ and it is standard to choose $\sigma = 1$ in calibrated macroeconomic models.

reform of reducing inefficiencies; and iii) that it leads to a drag on TFP growth. We want to emphasize that calibrating political effects properly is not easy. We will nevertheless make a rough attempt in order to get a sense of how the previously obtained welfare losses compare with the potential political costs of increased spending.

In a series of papers, Lane and Tornell analyze the "voracity effect" which they define as "a more than proportional increase in redistribution in response to an increase in the rate of return" (see Tornell and Lane, 1999; Lane and Tornell, 1996) . In their model, this is caused by wasteful investments, i.e., the equivalent of the κ's in our model. They present some suggestive empirical evidence of this waste (from case studies of Venezuela, Nigeria, and Trinidad and Tobago) and show that oil windfalls resulted in substantial increases, in what seems to be wasteful, government spending. To parsimoniously capture the fact that the voracity effect intensifies during oil windfalls, suppose that κ_g is affected by the increase in the oil spending as follows:

$$\kappa_{g,t} = \rho\kappa_{g,t-1} + f(\Delta OilSpending_t), \qquad \rho \in (0,1)$$

where $\Delta OilSpending_t$ is the change in oil spending (as a share of GDP). We assume that $f \geq 0$ and strictly positive only when $\Delta OilSpending_t$ and that it is strictly positive. Thus, we are assuming that the additional waste is a function of the increase in spending of oil revenues. Apart from the voracity effect, it also represents absorption constraints as discussed by van der Ploeg and Venables (2012). As the economy "gets used to" the high income, the absorption capacity increases which is captured by the waste-decay parameter $\rho < 1$.[36] The lag of κ_g in the expression captures the possible persistence of such an effect (say, due to instilled corruption networks in government following the windfall). Note that, in the SAYG scheme, the spending is $p_t \times oil_t$ while, in the fund scheme, this spending is the 4% return of the funds value. Hence, waste is increased in both scenarios.

One can consider several alternative assumptions regarding the functional form of f in the above expression for the voracity effect. Of the three countries discussed by Tornell and Lane (1999), Nigeria—as a multi-ethnic Sub-Saharan state—perhaps bears the closest resemblance to Uganda. In Nigeria, government spending increased during the period 1970 to 1981 from 22% to 36% as a share of GDP (Tornell and Lane, 1999). During this period, Nigeria's GDP did not increase at all while the resource income as share of GDP increased from five to 30%. We do not know how much of this increased spending was wasteful but, since GDP stood still, it is hard to see how an increase in government spending could be economically motivated. Thus, it seems plausible that most, if not all, of the increase was wasteful (through, e.g., construction projects awarded to politically connected contractors). These numbers then imply that for every percentage point increase of the resource share of GDP, κ_g increases by 1.6 percentage points, i.e., $f(\Delta OilSpending_t) = 1.6 \times \Delta OilSpending_t$. We have run the simulation assuming that this f-function will apply also in Uganda (but using the oil revenues in Uganda of course). The left panel of Figure 11 shows how κ_g evolves over time under three scenarios. The left-most curve (solid line) shows the baseline assumption where κ_g is not affected by oil spending. The middle and right-most curves display, respectively, κ_g for the fund and SAYG cases assuming the constant equals 1.6 as in Nigeria. Here, κ_g is increased both under the fund and under SAYG scheme. However, the shift is more pronounced in the SAYG scheme since there the economy is "shocked" by a rapid increase in oil revenue following the large increase in oil production in the early years. In the fund case, on the other hand,

36. We use $\rho = 0.95$ which is the same as the fall in κ in the baseline setup of the model.

Figure 11: The "Voracity Effect" Under the Fund and SAYG

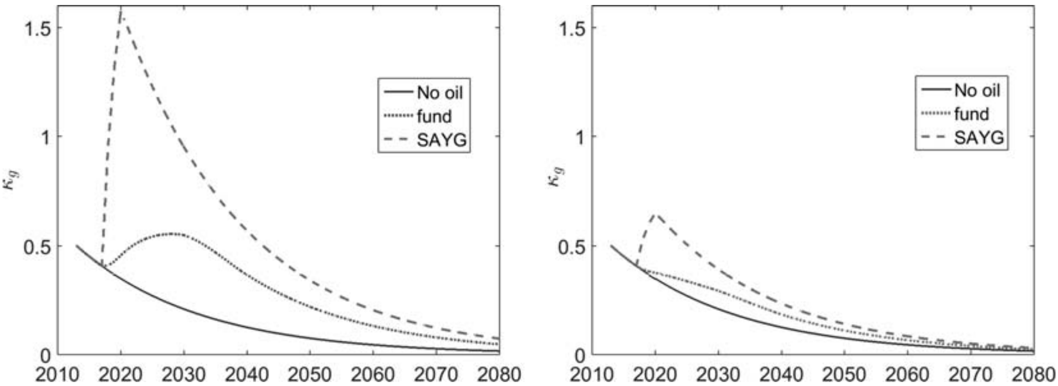

Left panel: Nigerian Experience, Constant = 1.6. Right panel: constant = 0.65. In both panels the black solid line depicts the base case without oil which is equivalent to the case with exogenous κ_g, i.e., without the voractiy effect.

these revenues are gradually phased in. Under this *f*-function, the fund scheme yields a higher welfare than the SAYG scheme. To see how robust this result is, we also analyzed how small the constant (which the Nigerian case suggested to be 1.6) has to be for welfare to equalize between the fund and the SAYG. The results show that the constant has to be below 0.65. That is, for the SAYG to be better, a one percentage point increase in oil spending cannot yield more than a 0.65 percentage point increase in waste—less than half of what the observations in Nigeria suggest. The evolution of κ_g in this case is depicted in the right panel of Figure 11.

As an alternative perspective of what oil-induced waste could imply, one can consider delayed reforms. Roughly speaking, this can be due to politicians buying support through increased spending rather than undertaking economic reforms to improve efficiency. As argued in Section 2.1, Uganda is in need of structural change. Suppose that the SAYG scheme implies a delay of reform so that the reduction in κ_g and κ is delayed by some time. Simulating different lengths of this delay, we find in the baseline setting of the model that a delay of one year is sufficient for the fund to be preferred over the SAYG scheme. Given the extensive empirical support for inefficiencies due to oil (see earlier discussions), a one year delay seems rather small.

Finally, delayed reforms also have the potential of harming TFP growth. As indicated in Section 2.2, the constant stream of income that can be achieved if all the oil could be taken out at once is equivalent to only 6 years of TFP growth. That the oil can be taken out at once is of course unrealistic, but that exercise is interesting as a perspective of the relation between the value of oil and productivity growth. To compare what a drag on productivity would imply for the welfare loss, we have made the following simple thought experiment. Suppose the SAYG, due to delayed or absence of reform, reduces TFP growth. How much would this *reduction* need to be in order for the welfare in the fund scheme to equal the welfare in the SAYG scheme in the baseline scenario? Our analysis suggests 0.006%. This shows that even marginal harm to growth quickly overshadows the gains of early usage of oil income.

It is of course hard to know the actual political effects increased spending in Uganda. However, to put our results about the drag in productivity growth into perspective, we can make some rough comparisons to previous findings in the resource-curse literature.[37] Sachs and Warner

37. It is important to note that different papers find different results with respect to the negative effect of resources on growth and that many find a conditional resource curse (see van der Ploeg 2011 for a discussion and Brunnschweiler and

(2001) estimate that a 1% increase in resource dependence is associated with a 0.05% decrease in annual growth over a thirty year period (see Table 2 ibid.). In Uganda, during the top years of the SAYG, oil will account for around 20–30% of GDP, which would imply, by the estimates of Sachs and Warner (2001), a decrease in annual growth of 1–1.5%. This is several orders of magnitude larger than the 0.006% we find to motivate the fund. Furthermore, the most serious political side effects of resources, such as civil conflict as suggested by Collier and Hoeffler (1998) and Garfinkel and Skaperdas (2007), imply welfare losses of several orders of magnitude larger (see, Abadie and Gardeazabal, 2003) compared to the difference between the fund and SAYG scheme.

We conclude from this exercise that if the fund and spending-constraint scheme can reduce the risks—of delayed reform, negative effects on growth or conflict—by even a very small amount, then that seems to more than compensate for the delayed access to the oil income. Getting the political constraints right—as one partly does by a clear and simple fund and spending rule—is more important than getting the spending profile right.

7. CONCLUDING DISCUSSION

In a developing economy, the merit of maximizing current spending from resource revenues is that it provides an opportunity to increase the consumption of the currently poor population. It also relaxes potential borrowing constraints and thereby enables increasing investment in domestic capital with high rates of return. Thus, constraining spending by utilizing a sovereign wealth fund with a restrictive spending rule may involve a welfare loss. Using Uganda as a case study, we analyze the size of this loss quantitatively and find, perhaps surprisingly, that it is most likely to be very small. Hence, given that having a sovereign wealth fund along with a simple, but rigid, spending constraint provides more transparency and helps decrease the potential negative political side effects associated with increased spending, adopting such a construction seems like a well motivated policy measure.

In light of the pronounced volatility of the oil price and the substantial evidence that government consumption in developing countries is pro-cyclical (Ilzetzki and Vegh, 2008), precautionary saving provides additional motivation for using a fund instead of spending revenues as they arrive (Cherif and Hasanov, 2013). One possibility is to use two funds: one for constraining spending over the long run and one with the specific purpose of smoothing government budgets in the short and medium run. Such a construction has, for instance, been used in Ghana (van der Ploeg et al., 2012) and could prove useful in Uganda as well. While focusing on the long-run costs and benefits of a sovereign wealth fund, we have abstracted from this possibility.

Our analysis has abstracted from other negative effects of resources, such as the potential for the Dutch Disease harming manufacturing exports or leading to increased income inequality. Incorporating these issues would favor the fund and spending-constraint construction as these negative effects would be delayed.

We also want to emphasize that the analysis we have performed is macroeconomic. That is, it has focused on the aggregate returns on domestic investments, which are indeed very high in Uganda and we found that, despite high returns, a fund scheme seems preferable. However, we have not looked at specific investments, so our conclusion does not mean that there do not exist projects in Uganda which have an even higher return or where the oil income would be handy to

Bulte 2008 for a paper challenging the existence of a resource curse). So our comparison here should be seen mainly as an exercise for what the ball-park size of potential effects may be.

use (for instance, we can imagine that schooling or investments in corruption would have high returns). Likewise, there might be projects with returns far lower than the aggregate rate. In an ideal world, returns to capital should equalize across projects. However, this is very unlikely to be the case in Uganda where efficient capital markets are missing and government project selection suffers from lack of political accountability. When evaluating the fund against the SAYG, one thus may wonder which of the schemes is most likely to lead to a more efficient prioritization of projects. Given the history of how resource wealth has led to wasteful investments, it seems that, as discussed in Section 6.3, this consideration speaks in favor of the fund.

The choice of Uganda for the quantitative analysis is primarily driven by our relative familiarity with the Ugandan context. Although the key assumptions we make regarding the long run development path are generally relevant to other developing countries as well, country-specific factors like the timing and quantity of the oil flow do influence our results. Thus, care should be taken in generalizing the results to a broader set of developing economies. Given the direct policy relevance of investigating the costs and benefits of a sovereign wealth fund and spending constraint, further research employing a similar analysis for other countries is well warranted to provide a more general picture on the relative merits of constraining spending.

APPENDIX

TFP Decomposition

The TFP decomposition uses the following model. Production is of the Cobb-Douglas type as represented in equation 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. The exponents γ and α are parameters assumed to be constant over time. We use data from the Penn World Table 7.1 for the period 1950 to 2010. Unfortunately, The Penn World Table does not contain data on capital stocks. However, it does have measures of gross investment, i.e., the sum of 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. Finally, assuming that the growth rates of public and private capital stocks are approximately equal, by differentiating the logarithm of equation (1) and rearranging, we can construct the growth rate of z_t, denoted g_{z_t}, by noting that

$$g_{z_t} = g_{y_t} - (\alpha + \gamma)(g_{k_t} - g_{l_t}) \tag{5}$$

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 set the depreciation rate to 10% per year and

38. The value of $\kappa = 1/3$ implies that to attain a certain increase in capital, Uganda has to invest 50% more than other countries. This seems like a reasonable assumption following discussions with Ugandan officials.

κ_t to $\frac{1}{3}$.[38] Finally, we set $\gamma = 1/6$ and $\alpha = 1/3$ which are standard numbers in the macroeconomic literature.

For now, we note that we can use (5) 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 (5), we obtain

$$g_y = \frac{g_z}{1-\alpha-\gamma}. \tag{6}$$

Under the previously made assumption that $\alpha = 1/3$ and $\gamma = 1/6$, GDP per capita thus grows twice as fast as productivity under balanced growth.

Parameter Values and Calibration Details

- $\alpha = 1/3$, representing a typical share of private capital of 1/3.
- $\gamma = 1/6$, representing a cost share of government capital of 1/6, although 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 $\rho_1 = 0.95$. Thus, we take the initial growth situation to be one where total factor productivity grows relatively fast, indicating a 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 assumptions is 2.5%, which is still somewhat 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 earlier. The convergence in productivity is rather slow, with a 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 as 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.

ACKNOWLEDGMENTS

We wish to thank the International Growth Center for financial and logistic support. Thanks also to Marguerite Duponchel for practical support in Uganda, and Richard Newfarmer, James Smith, Jakob Svensson and two anonymous referees for valuable comments. We are grateful to officials at the Ugandan Ministry of Finance, the Ministry of Energy and Mineral Development, the Central Bank of Uganda, and representatives of Tullow Oil, CNOOC and Total for sharing information.

REFERENCES

Abadie, A. and J. Gardeazabal (2003). "The economic costs of conflict: A case study of the basque country". *The American Economic Review*, 93(1): 113–132. http://dx.doi.org/10.1257/000282803321455188.

Ades, A. and R. Di Tella (1999). "Rents, competition, and corruption". *The American Economic Review*, 89(4): 982–993. http://dx.doi.org/10.1257/aer.89.4.982.

Aghion, P., P. Bacchetta, R.Ranciére, and K. Rogoff (2009). "Exchange rate volatility and productivity growth: The role of financial development". *Journal of Monetary Economics*, 56(4): 494–513. http://dx.doi.org/10.1016/j.jmoneco.2009.03.015.

Bacon, R. and S. Tordo (2006). "Experiences with oil funds: Institutional and financial aspects". Technical report, Energy Sector Management Assistance Program (ESMAP), World Bank.

Barnett, S. and R. Ossowski (2003). "Operational aspects of fiscal policy in oil-producing countries". In J. Davis, R. Ossowski, and A. Fedelino, *Fiscal Policy Formulation and Implementation in Oil-Producing Countries*, Washington, DC. IMF.

Berg, A., R. Portillo, S.-C. Yang, and L.-F. Zanna (2013). "Public investment in resource-abundant developing countries". *IMF Economic Review*, 61(1): 92–129. http://dx.doi.org/10.1057/imfer.2013.1.

Brunnschweiler, C. N. and E. H. Bulte (2008). "The resource curse revisited and revised: A tale of paradoxes and red herrings". *Journal of Environmental Economics and Management*, 55(3): 248–264. http://dx.doi.org/10.1016/j.jeem.2007.08.004.

Cass, D. (1965). "Optimum growth in an aggregative model of capital accumulation". *The Review of Economic Studies*, 32(3): 233–240. http://dx.doi.org/10.2307/2295827.

Cherif, R. and F. Hasanov (2013). "Oil exporters' dilemma: How much to save and how much to invest". *World Development*, 52: 120–131. http://dx.doi.org/10.1016/j.worlddev.2013.06.006.

Collier, P. and A. Hoeffler (1998). "On economic causes of civil war". *Oxford Economic Papers*, 50(4): 563–573. http://dx.doi.org/10.1093/oep/50.4.563.

Collier, P. and A. Hoeffler (2004). "Greed and grievance in civil war". *Oxford Economic Papers*, 56(4): 563–595. http://dx.doi.org/10.1093/oep/gpf064.

Dasgupta, P. and G. Heal (1974). "The optimal depletion of exhaustible resources". *The Review of Economic Studies*, 41: 3–28. http://dx.doi.org/10.2307/2296369.

Davis, J., R. Ossowski, J. A. Daniel, and S. Barnett (2003). "Stabilization and savings funds for nonrenewable resources: Experience and fiscal policy implications". In J. Davis, R. Ossowski, and A. Fedelino, *Fiscal Policy Formulation and Implementation in Oil-Producing Countries*, Washington, DC. IMF.

EDP Bill (2012). "Petroleum (Exploration, Development and Production) Bill No. 1" Ugandan Parliament.

Engel, E. and R. Valdés (2000). "Experiences with oil funds: Institutional and financial aspects". Technical Report 00/118, IMF Working Paper.

Garfinkel, M. R. and S. Skaperdas (2007). "Chapter 22 economics of conflict: An overview". In T. Sandler and K. Hartley, *Handbook of Defense Economics Defense in a Globalized World*, volume 2 of *Handbook of Defense Economics*, pages 649–709. Elsevier. http://dx.doi.org/10.1016/s1574-0013(06)02022-9.

Gelb, A. (1988). *"Oil Windfalls: Blessing or Curse?"*. Oxford University Press.

Gourinchas, P.O. and O. Jeanne (2006). "The Elusive Gains from International Financial Integration". *Review of Economic Studies*, 73(3): 715–741. http://dx.doi.org/10.1111/j.1467-937X.2006.00393.x.

Hamilton, J. D. (2009). “Understanding Crude Oil Prices”. *The Energy Journal*, 30(2): 179–206. http://dx.doi.org/10.5547/issn0195-6574-ej-vol30-no2-9.

Hart, R. and D. Spiro (2011). “The elephant in Hotelling’s room”. *Energy Policy*. 39(12): 7834–7838. http://dx.doi.org/10.1016/j.enpol.2011.09.029.

Henstridge, M. and J. Page (2012). “Managing a modest boom: Oil revenues in uganda”. OxCarre Research Paper 90, OxCarre (Oxford Centre for the Analysis of Resource Rich Economies).

Hotelling, H. (1931). “The economics of exhaustible resources”. *Journal of Political Economy*, 39(2): 137–175. http://dx.doi.org/10.1086/254195.

Hurlin, C. and F. Arestoff (2010). “Are Public Investments Efficient in Creating Capital Stocks in Developing Countries?”. *Economics Bulletin*, 30(4): 3177–3187.

Ilzetzki, E. and C. A. Vegh (2008). “Procyclical fiscal policy in developing countries: Truth or fiction?”. Working Paper 14191, National Bureau of Economic Research.

Isham, J. and D. Kaufmann (1999). “The forgotten rationale for policy reform: The productivity of investment projects”. *The Quarterly Journal of Economics*, 114(1):149–184. http://dx.doi.org/10.1162/003355399555972.

Jones, C. I. (2013). *“Macroeconomics”*. W. W. Norton & Company, third edition edition.

Koopmans, T. (1963). “On the concept of optimal economic growth”. Technical Report 163, Cowles Foundation Discussion paper.

Landon, S. and C. Smith (2015). “Rule-based resource revenue stabilization funds: A welfare comparicon”. *The Energy Journal*, 36(2): 117–43. http://dx.doi.org/10.5547/01956574.36.2.6.

Lane, P. R. and A. Tornell (1996). “Power, growth, and the voracity effect”. *Journal of Economic Growth*, 1(2): 213–241. http://dx.doi.org/10.1007/BF00138863.

Little, I. M. D., R. N. Cooper, W. M. Corden, and S. Rajapatirana (1993). *“Booms, crisis, and adjustment”*. Oxford University Press (for the World Bank).

Loayza, N. V., R. Ranciére, L. Servén, and J. Ventura (2007). “Macroeconomic Volatility and Welfare in Developing Countries: An Introduction”. *World Bank Economic Review*, 21(3): 343–357. http://dx.doi.org/10.1093/wber/lhm017.

Murphy, K. (1983). *“Macroproject development in the Third World”*. Westview Press.

Petroleum Intelligence Weekly (2015). *“East African Oil Gains Momentum, Pitfalls Remain”*. Energy Intelligence.

Pritchett, L. (2000). “The tyranny of concepts: Cudie (cumulated, depreciated, investment effort) is not capital”. *Journal of Economic Growth*, 5(4): 361–384. http://dx.doi.org/10.1023/A:1026551519329.

Sachs, J. D. and A. M. Warner (2001). “The curse of natural resources”. *European Economic Review*, 45:827–838. 15th Annual Congress of the European Economic Association. http://dx.doi.org/10.1016/s0014-2921(01)00125-8.

Sala-i-Martin, X. and A. Subramanian (2003). “Addressing the natural resource curse: An illustration from Nigeria”. Working Paper 9804, National Bureau of Economic Research.

Segura, A. (2006). “Management of oil wealth under the permanent income hypothesis: The case of São Tomé and Príncipe”. Technical Report Working Paper No. 06/183, IMF, Washington, DC:. http://dx.doi.org/10.5089/9781451864434.001.

Spiro, D. (2014). “Resource prices and planning horizons”. *Journal of Economic Dynamics and Control*, 48: 159–175. http://dx.doi.org/10.1016/j.jedc.2014.08.002.

Tornell, A. and P. R. Lane (1999). “The voracity effect”. *The American Economic Review*, 89(1): 22–46. http://dx.doi.org/10.1257/aer.89.1.22.

van der Ploeg, F. (2011). “Natural Resources: Curse or Blessing?”. *Journal of Economic Literature*, 49(2): 366–420. http://dx.doi.org/10.1257/jel.49.2.366.

van der Ploeg, F. and A. J. Venables (2011). “Harnessing windfall revenues: Optimal policies for resource-rich developing economies”. *The Economic Journal*, 121(551): 1–30. http://dx.doi.org/10.1111/j.1468-0297.2010.02411.x.

van der Ploeg, F. and A. Venables (2012). “Natural resource wealth: The challenge of managing a windfall”. *Annual Review of Economics*, 4: 315–337.

van der Ploeg, F., R. Stefanski, and S. Wills (2012). “Harnessing oil revenues in ghana”. The IGC Working Paper 12/0034.

Venables, A. J. (2015). “Depletion and development: natural resource supply with endogeneous field opening”. *Journal of Association of Environmental and Resource Economists*, forthcoming.

Vicente, P. C. (2010). “Does oil corrupt? evidence from a natural experiment in west africa”. *Journal of Development Economics*, 92(1): 28–38. http://dx.doi.org/10.1016/j.jdeveco.2009.01.005.

Young, A. (2012). “The african growth miracle”. *Journal of Political Economy*, 120(4): 696–739. http://dx.doi.org/10.1086/668501.