KAPSARC

مركز الملك عبدالله للدراسات والبحوث البترولية

King Abdullah Petroleum Studies and Research Center

Managing Oil Revenue Stabilization Funds: A Framework for Developing Policies

Nader AlKathiri, Tarek N. Atalla, Frederic H. Murphy and Axel Pierru

October 2017 KS-2017--DP021

About KAPSARC

The King Abdullah Petroleum Studies and Research Center (KAPSARC) is a non-profit global institution dedicated to independent research into energy economics, policy, technology and the environment, across all types of energy. KAPSARC's mandate is to advance the understanding of energy challenges and opportunities facing the world today and tomorrow, through unbiased, independent, and high-caliber research for the benefit of society. KAPSARC is located in Riyadh, Saudi Arabia.

Legal Notice

Key Points

Oil revenue stabilization funds provide short-run protection against oil revenue fluctuations – in the way that Saudi government deposits and reserve at the Saudi Arabian Monetary Authority (SAMA) have historically served as a buffer to decouple government budget from oil revenue fluctuations. By contrast, sovereign wealth funds create income for future generations to replace revenue streams from depletable resources – one of the purposes of Saudi Arabia's Public Investment Fund.

We developed a framework for optimizing policies for adding to and withdrawing from stabilization funds, which we apply to Saudi Arabia as a case study based on publicly available data. The quantitative results are sensitive to the specific assumptions on the likelihood of particular oil prices arising but the overall results are robust to a wide range of assumptions.

> In general, the results match intuition: Withdrawals from the fund occur when oil prices are low and these withdrawals are larger when the fund is larger. However, general intuition is not sufficient to capture the value of optimizing the policies.

> In our simulation, assuming that the Saudi government had applied the policies with the benefit of hindsight during the period 2003-2015, we find that the optimization approach can provide the same aggregate economic welfare during the period, but with a reserve fund that is $115 billion larger than the actual outcome at the end of the period.

> Looking forward, and assuming random paths for oil prices, the simulated policies lead to fund sizes that fluctuate over time with an average level of $14,512 per capita ($307 billion in total, using the 2015 Saudi population). There is less than a 5 percent probability that the fund exceeds $41,000 per capita ($866 billion in total) and the fund balance falls to zero 16 percent of the time during the simulation period.

> By dividing the fund into tranches, we estimate the probability distribution of the first time each tranche would be tapped (i.e., when the lower tranches have been fully used) and show that the higher tranches can be invested in longer duration assets that provide higher returns. This demonstrates why the boundary between a sovereign wealth fund and a stabilization fund is not a bright line.

Summary

Fluctuations in government expenditures result in welfare losses for risk-averse households and negatively impact the investment climate. Various oil-exporting countries have created revenue stabilization funds to cover the short- and medium-term expenses of government so as to reduce economic volatility. Even though Saudi Arabia has not officially established a stabilization fund, the government deposits and reserve at the Saudi Arabian Monetary Authority (SAMA) have historically served as a buffer to decouple government budget from oil revenue fluctuations.

The nature of the revenue movements determines the character of the fund. Due to economic forces, technological changes and new discoveries, oil revenues tend to stay within the same band and then switch to another band, which we term a regime. We use a Markov switching regime model to represent the movements among regimes. The model fits the random ways oil markets can switch from being oversupplied, balanced or tight. When estimating the model for per-capita oil revenues in Saudi Arabia, we find three different regimes.

Using these oil-revenue regimes, we develop a model that optimizes the buildup and draw down of a stabilization fund. The key factors we take into account beyond oil revenue movements are the effects of government spending on household well-being. The optimal policy consists of additions and withdrawals that are a function of the fund levels and oil revenues with the goal of improving household well-being. In practice, implementing the policy is easy since it depends on the actual revenues without regard to the regime.

The policies we develop satisfy common sense in that the larger the fund, the smaller the contribution during periods of high oil revenues and the bigger the withdrawal during periods of low oil revenues. Assuming that the Saudi government had used these policies during the period 2003-2015, we determine the evolution of the fund size by applying the policies to the historically observed oil revenue. The pattern of the optimal buildup is similar to that historically observed for the government deposits and reserve, but the differences illustrate the potential for improved economic welfare gains from an optimization approach.

Simulating random paths for oil revenues captures the potential fluctuations in the size of the stabilization fund. The fund level is below $41,000 per capita ($866 billion in total when using the 2015 Saudi population) 95 percent of the time and has an expected value of $14,512 per capita ($307 billion in total). The buildup of funds rarely achieves the maximum fund size before the market switches from a high to a low-revenue regime. The fund is empty 16 percent of the time because of extended stretches of low oil revenues. Starting from the average size, as a worst case it provides almost five years of supplementary revenues before reaching zero, during which the economy can adjust to a collapse in oil revenues.

Our approach also offers insight for the fund investment strategy. We divide the fund into tranches and consider that the lower tranches are tapped first. For each tranche we can estimate the probability distribution of the first time the tranche is tapped. The shape of the probability distribution is skewed to longer time periods, with the higher tranches having longer time distributions that are more stretched out. For instance, starting at a fund size of $42,000 per capita (≈ $900 billion in total) and breaking it down into six tranches, tapping the last tranche would occur at year six in the worst case and the average time to start tapping is 18.5 years. This means the higher tranches can be invested in longer duration assets that provide higher returns. These assets transition into the kinds of assets that are appropriate for a sovereign wealth fund.

Introduction

A country that depends on volatile export revenues from natural resources faces two questions related to these revenues: how should it prepare to replace the lost income stream as the resource depletes and how can it dampen the effects of revenue fluctuations on the economy and households when export revenues from the sale of the commodity are a major funding source? These questions have been addressed by creating funds called sovereign wealth funds and income stabilization funds. Some countries have just one fund and use it for both stabilization and preparing for resource depletion. Others maintain two separate funds, providing a firewall that protects the long-run funds from being used to meet short-run needs.

Bagattini (2011) finds in a cross-country analysis that separately identified stabilization funds lead to better economic outcomes because of improved governance. Good governance has been made more important by the lack of analytical methods that provide rigor to the decisions of how much to add or withdraw from a fund. In this paper, we begin the process of adding analytical rigor to managing these funds by using a model to optimize the buildup and drawdown of a stabilization fund.

Since both sovereign wealth funds and income stabilization funds are financial assets, they provide the best returns when their financial management is integrated. A country needs to know how much to add to financial assets during periods of high oil revenues, the amount that should be devoted to stabilization and how much should remain untouched to provide for subsequent generations. These are complicated, interrelated decisions.

Published reports on fund management practices indicate that the standard approach to managing these funds is to use rules of thumb that behave in ways that match peoples' intuition on how such a fund should operate. The reason for this approach is that rules of thumb are transparent and transparency contributes to good governance. In this paper we take a first cut at going beyond rules of thumb. If in the interest of transparency a country prefers to stay with simple rules of thumb, our results provide a benchmark for designing rules that provide similar welfare gains.

We develop analytical methods for determining the amount of money to add to the stabilization fund during times of high revenues, and to withdraw during periods of declining revenues. Separating the question of managing reserves devoted to income stabilization from managing a combined sovereign wealth and income stabilization fund simplifies the analysis while providing meaningful insights into fund management.

The stochastic structure of the revenue movements is central to determining the policy for building up and drawing down the stabilization fund. As part of our future research, we will look at integrated buildup and utilization strategies for the combined fund. The long-run goal is to take a fully integrated view, going beyond the size of the fund to include government debt and currency reserves that cover international transactions. Adding currency reserves to a combined economic stabilization and sovereign wealth fund sets a floor on the size of the fund so as to maintain the credibility of the currency after an extended period of low revenues.

The key factors we take into account in our approach are the effects of government spending on households and the properties of the commodity revenue fluctuations. The optimal policy consists of decisions that are a function of the fund level and commodity revenue. This policy differs from the decision rules proposed

in the literature or actually employed to manage funds (at least based on disclosed information) in that we maximize the expected intertemporal utility of households. These rules are generally chosen for their simplicity and transparency, guaranteeing fund additions when revenues are high and constraining withdrawals during low revenues so that the funds cannot be over used to address short-run considerations. The policy we develop here can either be used directly, for calibrating decision rules, if that is politically preferred, or as an assessment tool for examining existing policies.

In our analysis we recognize that the value from building a stabilization fund can be undercut by other aspects of fiscal policy, such as borrowing to cover budget deficits, which is why a fully integrated analysis is necessary. For example, Norway issued debt to cover budgetary shortfalls during the last economic downturn. Borrowing can be appropriate when revenues are low as long as borrowing is not a vehicle for boosting the stabilization fund when revenues are high. Furthermore, the cost of borrowing is much lower when lenders see a strong balance sheet, arguing for borrowing by using analytically-based plans before depleting reserves.

Literature Review

Managing economies with large depletable-resource sectors has been an important topic of economic research. Much of the literature has focused on the 'resource curse,' where economies wind up worse off from having large endowments of natural resources versus economies with no natural resources and provides governance prescriptions that ameliorate the problem. See Frankel (2012) and Ross (2015) for surveys on the resource curse.

A portion of the literature has focused on the negative consequences of uncertain revenue streams from the sale of resources. The uncertainty surrounding the economic growth rate has a social cost for risk-averse agents, usually measured as the loss of welfare that a representative agent is willing to incur to remove fluctuations in consumption. Though this cost might be negligible in certain economies, it is not the case for countries that rely on commodity export revenue, as shown by Pierru and Matar (2014) for Saudi Arabia. Furthermore, Aghion and Banerjee (2005) find that revenue volatility impedes economic growth by creating two problems. First, it adds uncertainty to the investment climate, dampening the level of investment. Second, governments too often see price increases as permanent and price decreases as temporary (see for instance Landon and Smith (2010)). The problem with this behavior is that government expenditures become pro-cyclical, aggravating the resource-revenue volatility (Erbil 2011), and make the investment climate even less inviting. El Anshasy and Bradley (2012) also find macro policies to be pro-cyclical, yet see a lag on the growth of expenditures that points to more fiscal sustainability. These papers illustrate the extensive literature on macro policy for resource-dependent economies. See also Barnett and Ossowski (2002), Sturm et al. (2009), Lopez-Murphy and Villafuerte (2010) and Frankel (2011).

Coping with price volatility has led to another set of literature for reducing expenditure volatility by smoothening available revenues to the government, e.g., Bagattini (2011), Landon and Smith (2013, 2015) and Villafuerte et al. (2008). Mexico pursues one strategy that provides near-term stability; it hedges the price of oil to stabilize oil revenues at a cost. This approach is recommended by Devlin and Titman (2004) and can make sense when planning the national budget for the year ahead. However, forward markets do not help in the low-price phase of a commodity cycle that lasts for several years. Another approach is to maintain an income stabilization fund that builds during periods of high revenue and supplies funds when revenues are low. This approach makes macro policy less pro-cyclical, dampens the effects of commodity cycles that can last for several years and allows for a less painful adjustment to an extended low-revenue regime.

Existing Rules for Managing Funds

The focus in countries that have established formal funds and in the work of the IMF and World Bank has been on institutional design to address the governance and fiscal-policy challenges of resource revenues. Thus, the literature is aimed at the central problems associated with resource endowments in countries where governments might not meet their fiduciary responsibilities because of short planning horizons. We have found no literature that looks at optimizing a stabilization fund. In fact, one paper states that it is impossible to develop optimal policies for a stabilization fund when prices are unknown (Devlin and Titman 2004).

All of the well-run funds described in the literature have one common feature: they operate under simple clear rules. Al-Moneef (2009) and Alsweilem et al. (2015) summarize the rules, organizational structures and finances of different countries' sovereign wealth and stabilization funds. The two most successful funds cited in the literature are run by Norway and Chile. Norway sends all of its oil revenues to a fund named the Government Pension Fund Global and returns 4 percent of the value of the fund to the national budget to cover a non-oil structural deficit, amounting to 9 percent of the government budget. This simple rule does not adjust the percentages using oil prices or the size of the fund and, instead, makes the oil revenues portion of the budget dependent on the level of the fund, which is more stable than oil prices.

Copper represents roughly half of Chile's export revenues. The country has a more elaborate approach to managing government revenues from copper exports than Norway does for oil. Two advisory committees of the Ministry of Finance project the long-run growth of its gross domestic product (GDP) and copper prices to estimate future government revenues. The budget is formulated using these long-run projections and any deviation in copper revenues from what they would be in the long-run projections adds to or subtracts from the sovereign wealth fund. A stabilization fund gets any excess revenues beyond the amount designated for the sovereign wealth fund. The government can move monies among the stabilization fund, sovereign wealth fund or budget at the discretion of the finance minister. Withdrawals are made from the stabilization fund and added to the budget when there is a structural imbalance that is greater than 1 percent, according to Fermandois (2011).

The need for sound management and clear rules is illustrated by the fund that the island of Nauru created. The island held large deposits of phosphate that are now depleted. The abandoned mines cover 80 percent of the land area of the island and the silt from mine runoff killed 40 percent of the island's marine life. Their sovereign wealth fund, the Nauru Royalties Trust, was mismanaged with investments that included a failed London stage musical written by one of its fund managers (Shenon 1995), and funds were withdrawn at too rapid a rate. The result is an island that has little economic value and a population that has gone from having one of the world's highest per capita incomes to impoverishment.

Given the fears about poor investment management and excessive spending, these funds have to be designed to ensure responsible management. Norway's keep-it-simple-and-clear strategy is a response to the management concerns. Other current or potential rules are as follows.

> Take interest-only payments from the fund and transfer any excess in the stabilization fund to a sovereign wealth fund that can be invested in longer-term assets.

- Endow ongoing budget activities:
  - General or public-employee pensions.
  - Higher education or education in general.
- Create an infrastructure-maintenance fund.
- Maintain a fund with no use restrictions but strict rules on additions and withdrawals:
  - Add and withdraw monies from the fund using moving averages of revenues and/or asset values.
  - Have a fixed percentage of resource revenues added and a different fixed percentage used for withdrawals.

Despite the well-documented problems associated with resource wealth, countries and state and provincial governments have generally been able to capitalize on the resource income streams to increase the well-being of their citizens. Given that some countries have escaped the resource curse and have disciplined approaches to manage their resource wealth, it is useful to go beyond the basics of institutions and constraining rules and ask what is the optimal policy for building up and drawing from the revenue stabilization fund? We focus on this question because countries that have successfully managed the resource curse can go beyond just avoiding pitfalls and better management of their resource endowments. Discretionary policies for managing stabilization funds often are initiated too late in the commodity cycle and rules can be too mechanical and insensitive to circumstances, according to Gylfason (2011). Even if countries prefer simple rules for transparency, benchmarking the rules against an optimal policy provides an estimate of the efficiency of their rules.

Framing the Fund Management Policy

We treat two aspects of the management policy for the stabilization fund. The first is the buildup and drawdown policy, and the second is how this policy interacts with the investment strategy. For the buildup and drawdown strategy, we draw upon the analytics of managing a strategic petroleum reserve. In oil-importing countries, a strategic petroleum reserve can lessen the economic effects of a surge in oil prices (Teisberg 1981; Hogan 1983; Murphy, Toman and Weiss 1986 and 1987).

The stabilization fund should have a significant share of its assets invested such that they can be drawn upon during periods of low revenues without incurring significant capital losses or requiring a different set of instruments. Although we do not address the mix of assets appropriate for a fund, we examine the maturity structure of the assets as a function of the potential timing of withdrawals in periods of low revenues and the asset levels in the fund. We show that a share of the stabilization fund can be invested in longer duration assets that provide higher returns. These assets transition into the kinds of assets that are appropriate for a sovereign wealth fund.

Managing Additions and Withdrawals to the Stabilization Fund

The nature of revenue changes of commodities determines the character of the stabilization fund. If the fluctuations were purely stationary around a known mean, then the fund would be more like normal cash reserves. However, commodities cycle through periods of high and low prices. This is due to a multitude of factors, including world demand and supply growth rates, new discoveries of the commodity, political events and technology change. A consequence of the possibility of extended periods of low revenues is that a stabilization fund is much larger than simple cash reserves and has to be managed differently.

This feature can be represented by a Markov switching regime model. In this model markets can switch among revenue levels in a standard Markov process with the switches in revenues differentiated from random fluctuations in revenues when within a regime. Given the number of independent events that impact oil markets, as mentioned above and detailed in Hamilton (2011), a Markov process fits the random ways oil markets can be oversupplied, in balance or tight.

The Markov model of switching regimes was developed by Hamilton (1989) for business cycle theory. It has been used in oil markets mainly to estimate volatility regimes in forward markets. Fong and See (2002) and Zhang and Zhang (2015) link regime changes in volatility to events. Deng (2000), and Higgs and Worthington (2008) estimate switching regimes for electricity prices. Pereira et al. (2016) examine 12 commodities – oil, gas and 10 metals – and cluster them into four different types of price and volatility regimes. Their analysis includes commodity super cycles, which is what we address here.

In this paper, we use a Markov process to model transitions among per-capita oil revenue regimes. We focus on revenues from commodity exports because both variations in prices and quantities exported impact revenue, and prices and quantities can be correlated, magnifying or dampening the effects of a regime change in prices. A per-capita approach accounts for population growth when designing rules for the stabilization fund.

To model the costs and benefits of the stabilization fund, we measure the effect of government spending on households. The government budget has three components, revenues from exports of the commodity, other revenues (e.g., from other exports and taxes paid by firms and households) and the net amount put into or withdrawn from the stabilization fund.

Let both i and j index the set of possible market states. The per capita revenue from the commodity of interest is R_i and the non-oil revenue is N_i. Let s be the per capita amount added to the stabilization fund during the period considered. A negative s represents a withdrawal from the fund. The per capita government budget for this period, $B_i(s)$, is calculated as follows:

$$B_i(s) = R_i + N_i - s \tag{1}$$

This budget is spent on public capital investment and other items lumped under 'government consumption,' denoted by g. These expenditures impact household consumption through different channels. As Linnemann (2006) emphasizes, the literature on the effect of government spending

on private consumption offers mixed views: a number of empirical analyses find that shocks to government spending seem to be associated with increases in private consumption, while the standard neoclassical intertemporal model is usually seen to predict the opposite. Thus, per capita household consumption is a function of the total budget $c_i=C_i(B_i(s))$. We assume that government consumption, g_i, is a function of the total government budget in state i, $g_i=G_i(B_i(s))$. In our illustrative case study of Saudi Arabia, we estimate these functions empirically.

The goods and services provided by government consumption add to household utility on top of the effects on the economy of the government expenditures. Following Barro (1981) and many others, we combine per capita household and government consumption into a single composite good that determines household utility. This reflects the benefit to households in the economy from both private and public goods. The composite good produced is given by $H(c_i(s),g_i(s))=H(C_i(B_i(s)),G_i(B_i(s)))$. The utility of the representative household is $U(H(c_i(s),g_i(s)))$. To simplify the notation, we introduce $V(i,s)\equiv U(H(c_i(s),g_i(s)))$.

Let f be the per capita amount in the stabilization fund at the start of the current period and let s be the amount added or removed. Then the amount in the fund becomes $f+s$. When including population growth, η, and assuming a return on the fund, r, the per capita fund size at the start of the subsequent period is

$$(f+s)\left(\frac{1+r}{1+\eta}\right).$$

When modeling total revenue, $\eta=0$ and the utility is simply the per capita utility multiplied by the population. Let $p(i,j)$ denote the probability of transitioning from state i to state j between two successive periods. Following the classical approach, the government maximizes welfare as measured by the household's current and future utility, calculated as the discounted sum of the utilities from each period, using the rate of time preference, ρ.

Because the fund balance in any year depends on actions taken in the previous year and the size of the fund at the start of that year, we have to represent this multi-period nature of the problem. A key feature of a stabilization fund is that it provides relatively short-run protection and does not address long-run issues such as resource depletion that are central for a sovereign wealth fund. Since we do not consider changes in revenues related to the resource depletion or other long-term structural shifts such as economic diversification, we approximate future welfare beyond the periods of interest using a steady-state representation of the economy. This means that when the model returns to, say, a regime with high oil revenues, the economy operates at the same baseline level when last in that regime. As a result, since the same welfare-maximization problem repeats in every period, the present value of expected utility $W(i,f)$, is a function of just the current state (i,f) and is independent of the time period. Thus, we do not need a time index. By letting r be the interest rate on the fund, the optimal policy is found by solving the following recursion:

$$W(i,f)=\max_{s\geq -f}\left[V(i,s)+\frac{1}{1+\rho}\sum_j p(i,j)\,W\left(j,(f+s)\left(\frac{1+r}{1+\eta}\right)\right)\right] \tag{2}$$

With:

$$s^*(i,f)=\arg\max_{s\geq -f}\left[V(i,s)+\frac{1}{1+\rho}\sum_j p(i,j)\,W\left(j,(f+s)\left(\frac{1+r}{1+\eta}\right)\right)\right]$$

That is, the expected utility in every state *(i,f)* is the maximum of the current household utility plus the present value of the expected future utility over possible fund additions or withdrawals. The inequality $s \geq -f$ in (2) ensures that the amount of money withdrawn cannot exceed the total in the fund. The recursion (1) is an infinite horizon dynamic program. Whenever a given market state and fund size recurs, the resulting current and expected future welfare is the same no matter the period you are in. Heuristically, this means we use a cyclic future rather than model a truncated future with simple growth-rate assumptions for the out years in the model.

The program (2) has a convenient reformulation that can be solved using linear programming (Winston 2004). Since *W(i,f)* is a maximum over all possible choices of s, we have

$$W(i,f) \geq V(i,s) + \frac{1}{1+\rho}\sum_j p(i,j)\, W\left(j, (f+s)\left(\frac{1+r}{1+\eta}\right)\right) \quad \forall s$$

At the optimal solution this expression is an equality for some s and that s is the optimal decision for *(i,f)*. Rearranging terms to have all variables on the left and constants on the right, we have the following linear program:

$$\min \sum_{i,f} W(i,f)$$

$$W(i,f) - \tfrac{1}{1+\rho}\textstyle\sum_j p(i,j)\, W\left(j, (f+s)\left(\tfrac{1+r}{1+\eta}\right)\right) \geq$$

$$V(i,s) \quad \forall s \geq -f \tag{3}$$

where the objective function ensures that for some s (3) is an equality for each *(i,f)*. To solve (3) as a linear program, we need to discretize the fund size and restate the decisions so that the decisions lead to a transition from the current fund size into another fund size in the discretization. The simplest approach is to apply uniform increments and decrements in the fund size and to have decisions that correspond to those increments and decrements. That is, instead of having a transition to a fund size of $(f+s)\left(\frac{1+r}{1+\eta}\right)$ from f, we want the transition to be to $f+u$, where u is a uniform discretization. That is,

$$u = (f+s)\left(\frac{1+r}{1+\eta}\right) - f = \frac{(r-\eta)f + (1+r)s}{1+\eta} \tag{4}$$

with u defined to make the discretization of the fund size uniform. Then s takes on the value

$$s = \frac{(1+\eta)u - (r-\eta)f}{1+r}$$

The decision u represents the change in the fund size between the start of the current period and the start of the following period, whereas the decision s represents the amount withdrawn from the fund during the period considered.

By substituting u, (3) becomes the following linear program:

$$\min \sum_{i,f} W(i,f)$$

Subject to:

$$W(i,f) - \tfrac{1}{1+\rho}\textstyle\sum_j p(i,j)\, W(j, f+u) \geq$$

$$V\left(i, \tfrac{(1+\eta)u-(r-\eta)f}{1+r}\right) \quad \forall u \geq -f \tag{5}$$

which can be readily discretized in u. In each state, $u^*(i,f)$, the value of u that maximizes $W(i,f)$, is obtained when (5) is binding. We can determine $W(.,.)$ and $u^{\wedge *}(.,.)$ for all states simultaneously.

From (4), the optimal policy for withdrawal or addition is $s^*(i,f)$ such that

$$s^*(i,f) = \frac{(1+\eta)u^*(i,f)-(r-\eta)f}{1+r} \tag{6}$$

To implement our approach we need to calibrate the functions $Bi(\bullet), Gi(\bullet), Ci(\bullet), H(\bullet,\bullet), U(\bullet)$, solve (5) as a linear program, and derive $s^*(i,f)$. We present the results for Saudi Arabia in the next section.

Calibration for Saudi Arabia

The numerical results are dependent on the assumptions used in calibrating the model, including oil revenue regimes, interest rate, household risk aversion or the preference parameters that govern the interactions between private and government consumption in the utility functions of households. The chosen quantitative values of any of these parameters can lead to significantly different specific numerical outcomes. However, sensitivity analyses indicate that the qualitative characteristics of the results hold across scenarios. We are not advocating that a particular set of assumptions should be used. Instead, we view this work only as a proof of concept for managers of a stabilization fund who will probably need to make assumptions that are very different from the ones we use here when actually planning a fund.

Estimating a Markov switching regime for per capita oil revenues

To construct data series on a quarterly basis, we first calculate the product of the monthly Saudi crude oil production (source: IEA) and average oil price (source: EIA). The resulting monthly values are combined to produce quarterly estimates of the value of the oil produced. Each year's quarterly estimates are then rescaled by the same coefficient so that their sum matches the Saudi Arabia's annual oil revenue reported for that year by the Saudi Arabian Monetary Authority (SAMA, 2016). We use the resulting rescaled values as quarterly oil revenues for 1984-2015, providing 128 data points. The quarterly data are deflated using the World CPI (source: St Louis Fed 2016), and divided by the Saudi population in each year (source: CDSI). All data in this paper are expressed in U.S. dollars of the year 2015.

The model is estimated using Eviews 9 following the approach from Filardo (1994), Kim and Nelson (1999), Benz and Truck (2006) and Sims and Zha (2006). The specification is based on a three-stage model that takes the logarithm of the quarterly per capita oil revenue as the dependent variable and a constant as the switching variable. In each regime the per capita revenue is therefore assumed to follow a lognormal law, the expected value of which is given by Table 1. For each regime, the null hypothesis of a normal distribution of errors is not rejected by the Jarque-Bera test at a 5 percent level.

Table 1. Estimated oil revenue regimes over 1984-2015.

Regime specifications	Revenue level	LOG coefficient	Std. error	Z-statistic
Per capita oil revenue regimes				
Collapse Regime	$747 per capita	6.59	0.09	69.54
Low Regime	$1,550 per capita	7.32	0.07	166.06
High Regime	$3,415 per capita	8.11	0.06	135.39

Source: KAPSARC.

The estimation results in a specification with three possible per capita quarterly oil revenue regimes: Revenue-collapse, Low and High.

Tables 1 and 2 show the estimates of the regime levels and the associated Markov probability transition matrix, respectively.

The dominant feature of the Markov matrix in Table 2 is that once in a state, the probability of remaining in that state is high. That is, the regimes are persistent. The steady-state probabilities are the long-run probabilities of being in a regime at any point in time. This means that there is only a 5 percent chance of being in a revenue-collapse regime and the most likely regime is low.

Figure 1 shows the regime levels and switching among regimes estimated over the sample period for per capita revenues.

Estimating non-oil revenues

We assume that non-oil revenues are proportional to oil revenues for each revenue regime:

$$N_i = \pi_i R_i \tag{7}$$

We estimate π_i using the following equation:

$$\frac{non-oil\ revenue}{oil\ revenue} = \alpha_1 + \alpha_2 D_L + \alpha_3 D_H \tag{8}$$

Where D_L and D_H are dummy variables for low and high oil revenue regimes respectively, with: $\pi_C = \alpha_1 = 0.46$, $\pi_L = \alpha_1 + \alpha_2 = 0.29$, and $\pi_H = \alpha_1 + \alpha_3 = 0.14$.

The results are shown in Table 3.

Table 2. Probabilities estimated for the Markov matrix.

Markov probability matrix of transitions among regimes

	Per capita oil revenue regimes		
	C	L	H
Collapse regime (C)	83.9%	16.1%	0.0%
Low regime (L)	1.4%	93.4%	5.2%
High regime (H)	0.0%	6.9%	93.1%
Steady state probability	5%	54%	41%

Source: KAPSARC.

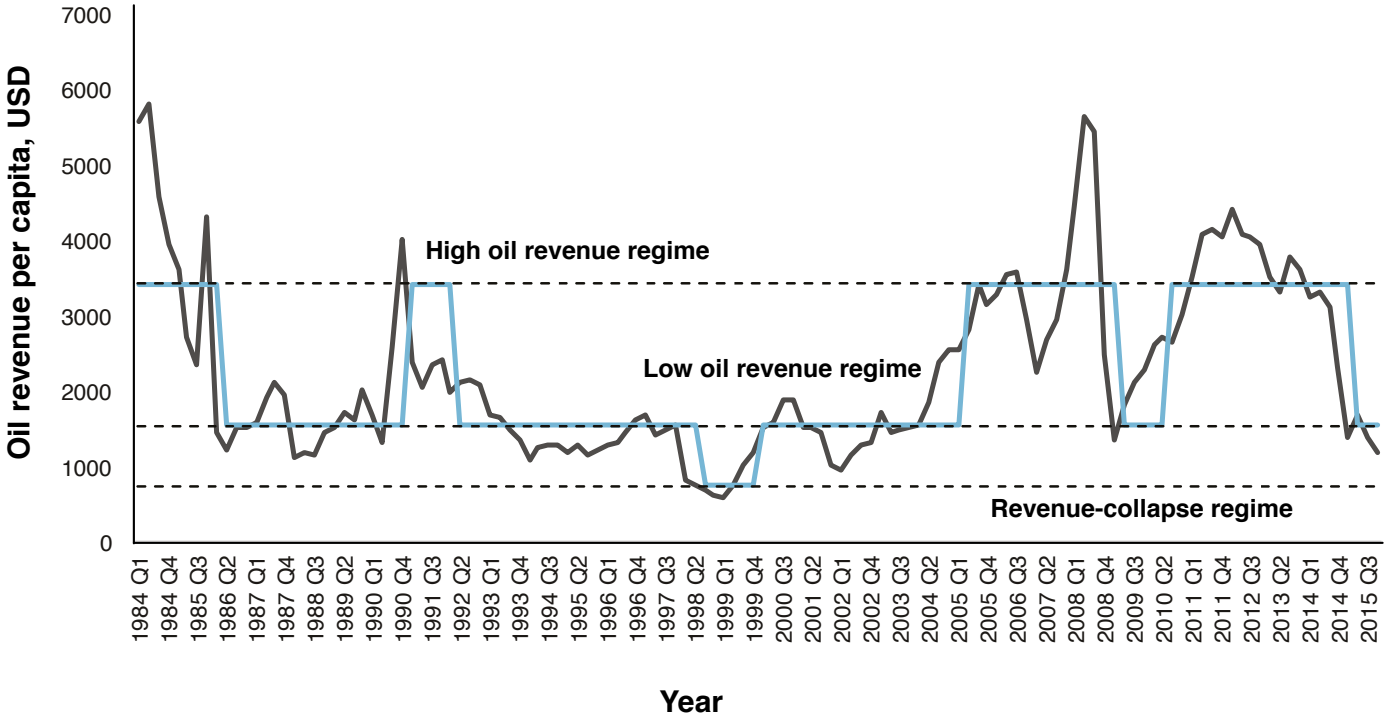

Figure 1. Estimated regimes plotted against actual oil revenues.

Source: KAPSARC.

Table 3. Estimated parameters for the ratio of non-oil revenue to oil revenue.

Parameters	Value of the coefficient	Std. error	t-statistic	p-value
α_1	0.46	0.04	10.93	0.00
α_2	-0.17	0.07	-2.48	0.02
α_3	-0.31	0.06	-4.94	0.00

Source: KAPSARC.

Calibrating the government consumption function

We define government consumption to be total government expenditures less capital expenditures. Let φ_i be the ratio of the government consumption to the total budget spent, which we assume to be constant within each regime. Then government consumption is defined as

$$G_i(B_i(s)) = \varphi_i B_i(s) \tag{9}$$

With dummy variables D_L and D_H for the medium and high revenue regimes, we estimate φ_i using the following regression:

$$\frac{Government\ consumption}{Total\ government\ expenditure} = \beta_1 + \beta_2 D_L + \beta_3 D_H \tag{10}$$

Based on the results in Table 4, we find:

$$\varphi_C = 0.87, \varphi_L = 0.73, \varphi_H = 0.71$$

Calibrating the private consumption function

We assume that per capita household consumption is a linear function of per capita government expenditure:

$$C_i(B_i(s)) = \omega + \tau B_i(s) \tag{11}$$

We use the KAPSARC Global Econometric Macroeconomic Model that is based on the Oxford Economics Global Economic Model (http://www.oxfordeconomics.com/) but includes an expanded representation of Saudi Arabia's economy. We use the reference scenario and build five alternative scenarios that consist of increasing annual government expenditures in 2012 by 10 percent, 20 percent, 30 percent, 50 percent and 80 percent, respectively. In each scenario, the model gives the resulting annual household consumption for 2012. We divide public and private consumption by four to express them on a quarterly basis. We estimate the coefficients by running a regression on these six data points.

Table 4. Estimated parameters for the ratio of government consumption to total expenditure.

Parameters	Value of the coefficient	Std. error	t-statistic	p-value
β_1	0.87	0.03	27.73	0.00
β_2	-0.14	0.05	-2.79	0.01
β_3	-0.16	0.05	-3.30	0.00

Source: KAPSARC.

Table 5 shows the estimated parameters. When the government increases its budget by $1,000 per capita, contemporaneous private consumption increases by $5 per capita.

Calibrating the composite good function

Here we assume that private and public consumption are perfect substitutes (see for instance Barro (1981)):

$$H(c,g)=c+g \tag{12}$$

The next section shows the optimal buildup and withdrawal policy under this assumption.

Utility function

Following Pierru and Matar (2014), who consider a coefficient of relative risk aversion of 2 for Saudi Arabia, we use the following constant-relative-risk-aversion utility function:

$$U(h) = -\frac{1}{h} \tag{13}$$

The representative household's utility is:

$$V(i,s) \equiv -\frac{1}{\omega+(\tau+\varphi_i)((1+\pi_i)R_i-s)}$$

We set the annual rate of time preference to 1 percent ($\rho = 0.25\%$) and choose an annual real interest rate of zero percent ($r = 0$). The annual population growth rate is set to 1.7 percent ($\eta = 0.425\%$), which is the estimated Saudi population growth rate for the period 2015-2025 (source: CDSI).

Table 5. Estimated parameters for the effect of government expenditures on private consumption.

Parameters	Coefficient estimate	Std. error	t-statistic	p-value
		Per-capita oil revenue		
ω	1,379.4	0.375	14,731.45	0.000
τ	0.005	0.000	225.55	0.000

Source: KAPSARC.

Results for the Saudi Arabia Case Study

An optimal policy for the buildup and drawdown of the stabilization fund consists of an amount to be added or withdrawn for each possible level of the fund and each oil revenue regime. In this section and the next we highlight features of the optimal policies, the behaviors of the funds and the implications of those behaviors for fund management. Again, we take the Markov matrix of per capita revenue as an estimate of the switching regime for illustration since our goal is to demonstrate that our approach offers insights into managing the stabilization component of a fund, not specific numbers at this stage of model development.

The resulting optimal policy for additions and withdrawals to the fund is given by Figure 2. When the policy provides amounts expressed on a per-capita basis, the numbers in parentheses indicate the corresponding absolute amounts, which is obtained by multiplying the per-capita amount by the Saudi population (21.13 million in 2015).

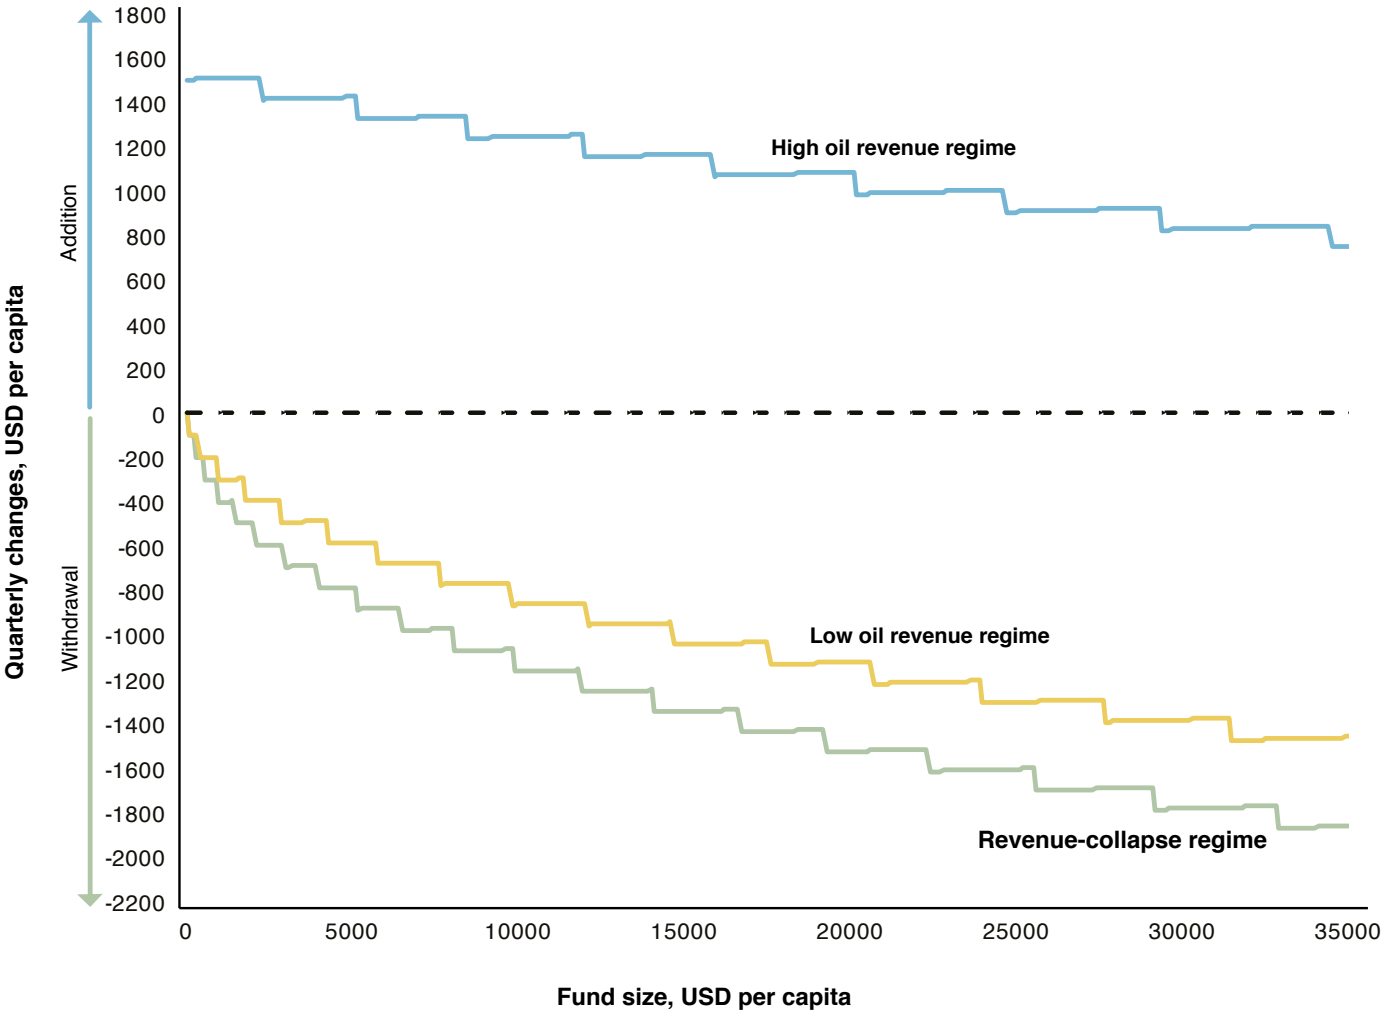

Figure 2. Optimal policies for per capita oil revenue regimes.

Source: KAPSARC.

In the high oil revenue regime money is added and the fund is increased up to a maximum of $89,100 per capita. In the low and revenue-collapse regimes, money is withdrawn until the fund is depleted or the high oil revenue regime returns. As expected, the larger the fund, the smaller the addition and the larger the withdrawal. In the low oil revenue regime, money is always withdrawn but at a lower rate than with the revenue-collapse regime. As an example, assume that, at the start of the quarter, the fund size is $7,500 per capita ($158.5 billion). During the quarter, the optimal policy is then to add $1,337 per capita ($28.6 billion) to the fund in the high oil revenue regime, to withdraw $671 per capita ($14.2 billion) in the low oil revenue regime and to withdraw $972 per capita ($20.5 billion) in the revenue-collapse regime.

How much protection the fund provides depends on its size at the start of a low or collapse oil revenue regime after building up during a regime with high oil revenues. We calculated the probability distribution of the fund size using steady-state probability distribution computed as in Winston (2004, 935), which gives the long-run probability of being in a state. Based on the optimal policy and the transition matrix estimated for oil revenue regimes, we built an extended transition matrix T for the bi-dimensional state (fund size and revenue regime) and determined the steady-state joint probability vector S such that: S=ST. We then derived the marginal fund-size probability distribution by adding the joint steady-state probabilities over the oil revenue regimes for each fund size.

Figure 3 shows the long-run probability distribution of the fund sizes. The size with the highest probability, 16 percent, is associated with a depleted fund, which is consistent with the purpose of a buffer that is created for short-run economic stabilization. The probability that the fund size exceeds $41,000 per capita ($866 billion) is only 5 percent, since it is unlikely that the fund can build beyond this level before oil revenues switch back to the low or collapse regimes and withdrawals start.

Note that the non-monotonicity of the curve is a feature of the optimal policy. As the fund is empty 16 percent of the time, the additions during the high revenue regime have high odds of starting from 0. Since the optimal policy has discrete additions determined by the fund size, when starting from 0, the buildup will follow exactly the same path and avoid other paths until the regime changes. The reason the curve is not even bumpier is because the revenue regime can change and the buildup does not always start at 0.

We gain two insights from examining the minimum time to run out. First, the length of time until the last dollar in the fund is used is quite long even in the worst case. Thus, as the stabilization fund grows, it blends smoothly into the sovereign wealth portion, where there should be no runout time, and there is no clear boundary between the two. Second, since assets with longer durations provide higher returns, we can begin to think about an investment strategy that divides the fund into tranches with each tranche invested taking into account the likelihood the funds would be needed after a given timespan.

The choice of investments for each tranche depends not only on the shortest time to being tapped but also on the probability distribution of the time until the tranche is tapped. In Figure 4, we divide the fund into six tranches. By simulating oil revenue regimes and applying the optimal policy, for each tranche we estimate the probability

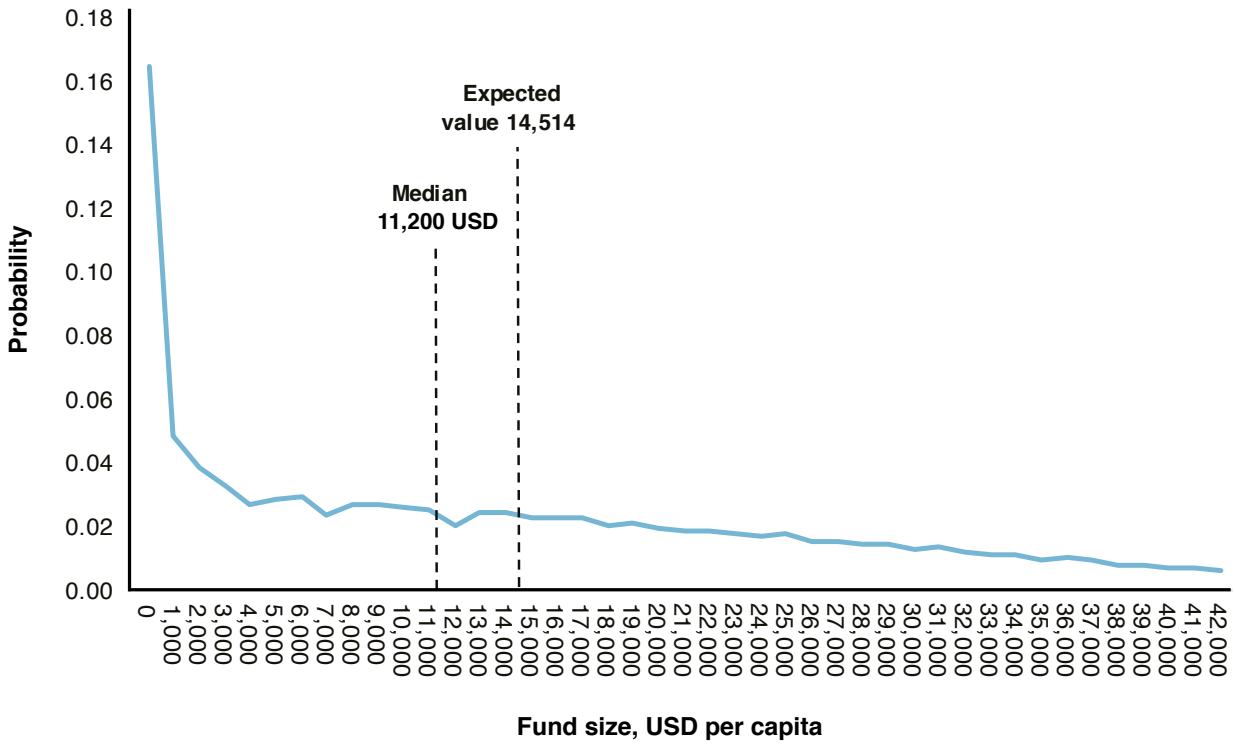

Figure 3. Probability distribution of the fund size.

Source: KAPSARC.

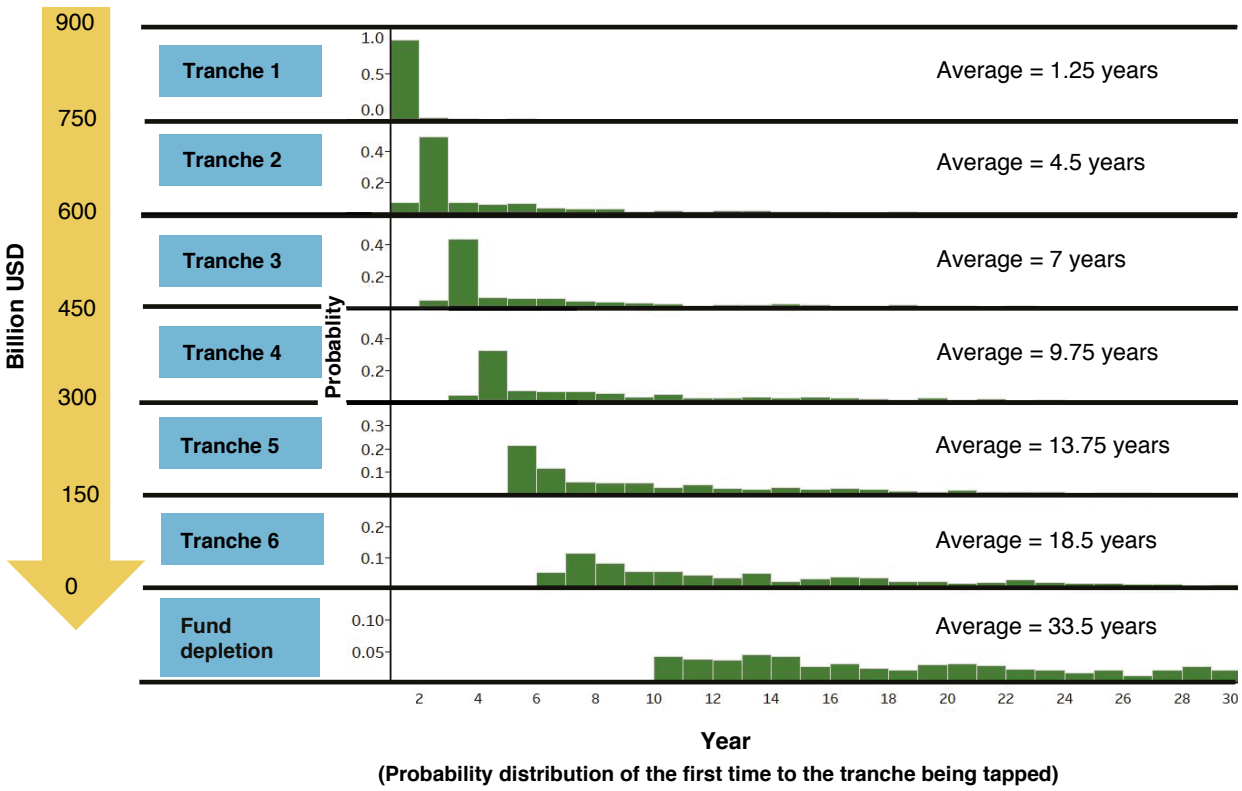

Figure 4. Probability distribution of the first time to the tranche being tapped.

Source: KAPSARC.

distribution of the first time to the tranche being tapped (i.e., when the lower tranches have been fully used). We also indicate the mean of each time distribution. We assume an initial fund of $900 billion (≈$42,000 per capita).

Note that we use relatively coarse tranches and a policy should consider an investment strategy that uses either smaller tranches or a continuous transition to longer-term assets. Nevertheless, Figure 4 provides the starting point for establishing an asset mix. A key feature in moving to higher tranches is that the distribution of time shifts further out at a faster rate and the distribution has an increasingly wider spread. This is consistent with the transition to a sovereign wealth fund that is managed like an endowment fund with long-lived assets, (Swensen 2009).

Comparing the Fund Policy with Actual Experience

All of the above results and analyses are done with three estimated revenue regimes while actual revenues can take a continuum of values. To make a policy practical, we need to convert the decisions associated with the specific values associated with revenue regimes to decisions for all possible revenue amounts. The simplest way to do that is to interpolate when the oil revenue is between two revenue regimes. That is, with a revenue R and $R_L \leq R \leq R_H$, let $R=\alpha R_L+(1-\alpha) R_H$ and $0 \leq \alpha \leq 1$, we add to the fund:

$$s = \alpha \times s(L,f) + (1-\alpha) \times s(H,f) \tag{14}$$

When the revenue is above R_H, we set the amount to:

$$s = s(H,f) + \frac{s(H,f)-s(L,f)}{R_H-R_L}(R-R_H) \tag{15}$$

This formula increases the additions beyond $s(H,f)$ by the rate of change of the addition when going from the low to high revenue regime multiplied by the increase in revenue beyond R_H. When the revenue is below R_C, the amount withdrawn is:

$$s = s(C,f) + \frac{s(L,f)-s(C,f)}{R_L-R_C}(R-R_C) \tag{16}$$

One aspect of comparing the actual experience with our model results is the extent to which Saudi Arabia has been able to keep to its annual budgets. The reason this is important is because budgeted savings were not always realized due to spending beyond the budget. Figure 5 shows the budgeted government expenditures and compares them with the actual expenditures. As can be seen, the actual expenditures tend to exceed what had been budgeted at the beginning of each year.

We now apply the optimal policy to the historically observed oil-revenue and obtain what we term a 'counterfactual fund size' series. In other words, we assume that the Saudi government had applied the policy shown in Figure 6 since the first quarter of 2003. For the initial fund size at the beginning of 2003, we use $617 per capita (SAMA 2017). This amount combines the government current account, the government reserve and allocated deposits for government projects in SAMA's balance sheet starting in 2003. Figure 6 shows the level of government deposits and reserves observed during the entire period being considered along with the counterfactual fund size. The figure also shows the potential fund size if the announced government budget had been realized with no deviation as illustrated in Figure 5 and we call it 'rigid budget discipline' series.

The timing of the buildups and drawdowns in the simulations matches the increases and decreases in actual funds. The counterfactual fund size tracks closely with the historical levels of government deposits and reserves until 2011, when we begin to see a higher accumulation of reserves using the optimal rule compared with actual contributions.

The welfare gains from using the optimal policy compared to the historical policy can be measured as the increase in the historical size of the fund at the end of the period that would equate the total welfare for both policies. That is, we need to increase the fund size in 2015 for the historical policy by the amount x such that:

$$\sum_{t=t_0}^{t=T}(1+\rho)^{T-t}\,V\big(i_t, s^*(i_t, f_t^*)\big) + W(i_T, f_T^*) =$$

$$\sum_{t=t_0}^{t=T}(1+\rho)^{T-t}\,U(c_t, g_t) + W(i_T, f_T + x)$$

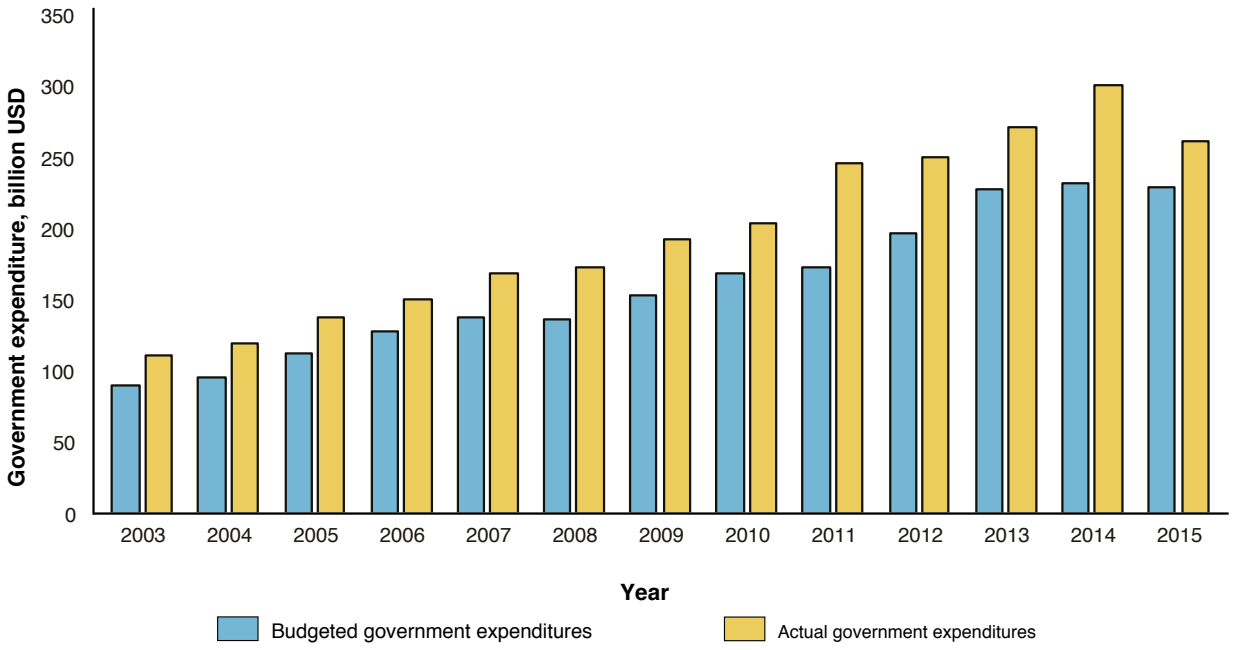

Figure 5. Budgeted government expenditures versus actual.

Source: KAPSARC.

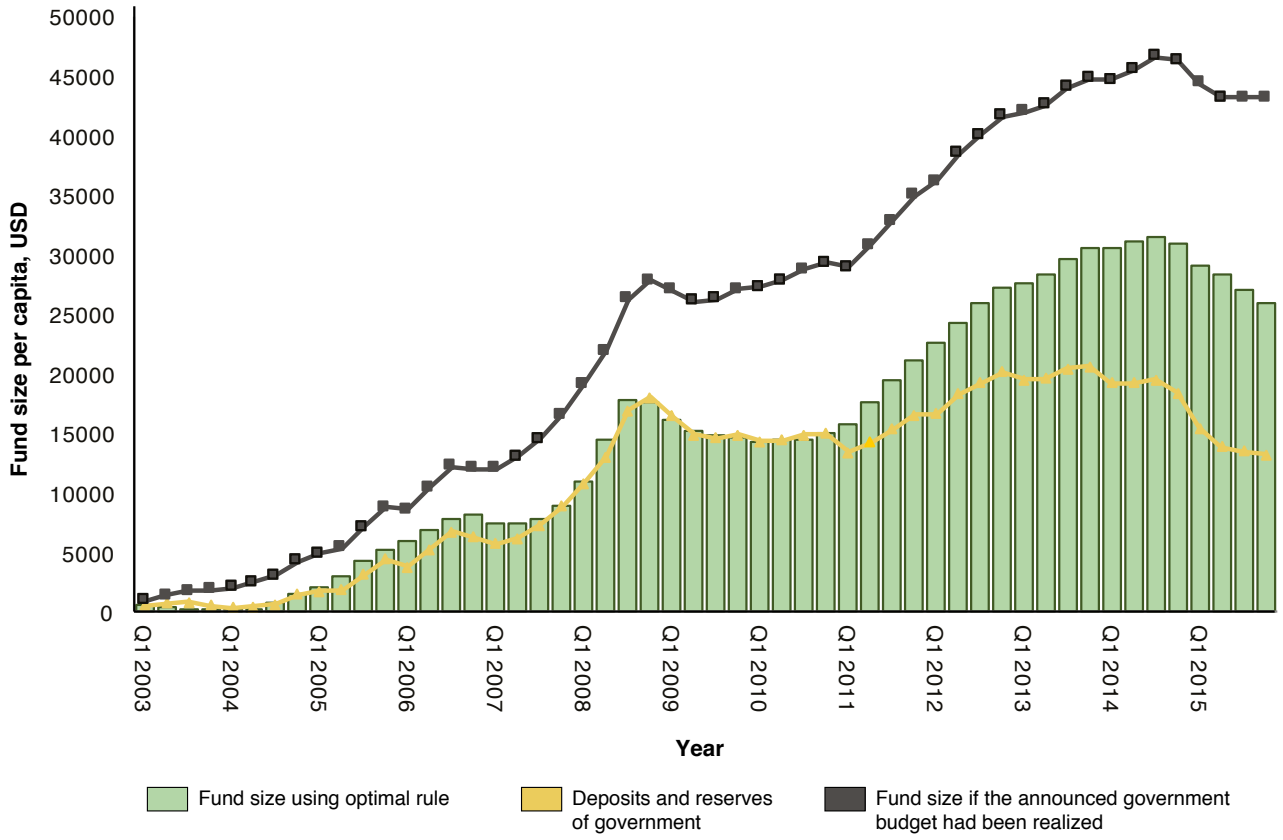

Figure 6. Historical and counterfactual fund size.

Source: KAPSARC.

Where c_t, g_t, and f_T are historical per capita private consumption, government consumption and fund size respectively. The fund size from the optimal policy is denoted by f_t*. The period covered is from t_0=*Q1* 2003 to *T*=*Q4* 2015.

The result shows that, with the benefit of hindsight, applying the optimal policy to achieve the same welfare outcome would have resulted in the balance of the stabilization fund being $115 billion higher at the end of the period compared with the historical policy. On the other hand, with the same approach, applying a rigid budget discipline would have resulted in a welfare loss equivalent to $238 billion compared with the optimal rule, despite building a higher fund at the end of the period.

Conclusions

Our analysis is the first effort to develop optimal policies for the revenue stabilization portion of a national fund. We present a new, analytical-based approach that can assist in managing stabilization funds of commodity exporting countries.

In the case of Saudi Arabia, the resulting policies roughly track with the recent history of fund increases and decreases. Furthermore, the policy has the intuitively desirable properties of a quicker buildup when the fund is low and greater withdrawals when the fund is large. This property is a provable feature of the results and follows from the decreasing marginal returns on expenditures in each state. It is not forced on the model.

Several aspects of the current model deserve further exploration before it is ready to assist developing a management policy. Most importantly, since the market has had relatively few transitions among regimes, the choice of probabilities and expected prices in the regimes should include judgments that compensate for the lack of data.

We do not account for habit formation and that the utilities of households can depend on both current and past consumption. Furthermore, we do not capture either the costs of adjusting public capital expenditures because of volatile budgets, or the loss of private investment due to greater risks in a more volatile economy. Including the lagged effects requires adding a third dimension to the state space in our analysis, the prior level of expenditures. Doing this, however, has a multiplicative effect on the computation time.

In addition, a government can finance spending with debt at an interest rate that increases with the country's debt level. Using debt as a stabilization tool can be handled in our model by extending the state space for the fund size to negative values.

Lastly, as we have mentioned throughout the paper, an integrated view of the financial structure of all funds, including sovereign wealth and currency stabilization funds, is the ultimate goal of this research project. Included in this model is not only the buildup and drawdown strategies but also allocations to the different funds and portfolio allocations to different asset classes based on the timing and odds of potential demands on the different tranches within funds as well as across funds.

The model can also be used for benchmarking existing rules and improving upon ad hoc additions to and expenditures from a fund. We expect to expand the utility of this model and develop a suite of models that cover all aspects of fund management.

A surprise to us is how long the short-run is when designing the stabilization portion of the fund. In retrospect this makes sense because price regimes can last for years. A consequence of potentially long time spans is that an optimally designed stabilization fund can reach large amounts and it is not like a reserve that covers day-to-day fluctuations in cash flows. Since the boundary with the sovereign wealth portion is fuzzy, our next step is to model an integrated stabilization and sovereign wealth fund.

References

Aghion, Philippe and Abhijit Banerjee. 2005. "Volatility and Growth: Clarendon Lectures in Economics."

Al-Moneef, Majid. (2009). “The Role of the Sovereign Wealth Funds in the Management of Petroleum Surpluses.” Arab Economic Journal. No. 47

Alsweilem, Khalid A., Angela Cummine, Malan Rietveld and Katherine Tweedie. 2015. "A comparative study of sovereign investor models: Sovereign fund profiles." Belfer Center for Science and International Affairs and Center for International Development: Harvard Kennedy School.

Alsweilem, Khalid A. 2015. "A Stable and Efficient Fiscal Framework for Saudi Arabia: The Role of Sovereign Funds in Decoupling Spending from Oil Revenue and Creating a Permanent Source of Income." Belfer Center for Science and International Affairs and Center for International Development: Harvard Kennedy School.

Bagattini, Gustavo Yudi. 2011. "The Political Economy of Stabilisation Funds: Measuring their Success in Resource-Dependent Countries." IDS Working Papers, no. 356: 01-60.

Barnett, Mr Steven and Mr Rolando Ossowski. 2002. Operational aspects of fiscal policy in oil-producing countries. No. 2-177. International Monetary Fund.

Barro, Robert J. 1981 "Output effects of government purchases." Journal of political Economy 89, No. 6: 1086-1121.

Benz, Eva and Stefan Trück. 2009. "Modeling the price dynamics of CO2 emission allowances." Energy Economics 31, No. 1: 4-15.

Centeral Department of Statistics and Information CDSI. 2010. "Mid-year population estimates for administrative regions and provinces." (Translated from Arabic.)

Davis, Mr Jeffrey M., Ms Annalisa Fedelino and Mr Rolando Ossowski. 2003. Fiscal policy formulation and implementation in oil-producing countries. International Monetary Fund.

Deng, Shijie. 2000. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Berkeley: University of California Energy Institute.

Devlin, Julia and Sheridan Titman. 2004 "Managing oil price risk in developing countries." The World Bank Research Observer 19, No. 1: 119-139.

El Anshasy, Amany A., and Michael D. Bradley. 2012 "Oil prices and the fiscal policy response in oil-exporting countries." Journal of Policy Modeling 34, No. 5: 605-620.

Erbil, Nese. 2011. "Is fiscal policy procyclical in developing oil-producing countries?"

Fermandois, Arturo. 2011. "Chile and its Earthquake, preparedness, response and lessons", http://dels.nas.edu/resources/static-assets/materials-based-on-reports/presentations/AmbassadorFermandois.pdf

Filardo, Andrew J. 1994 "Business-cycle phases and their transitional dynamics." Journal of Business & Economic Statistics 12, No. 3: 299-308.

Fong, Wai Mun and Kim Hock See. 2002 "A Markov switching model of the conditional volatility of crude oil futures prices." Energy Economics 24, No. 1: 71-95.

Frankel, Jeffrey A. 2001. A solution to fiscal procyclicality: The structural budget institutions pioneered by Chile. No. w16945. National Bureau of Economic Research.

Frankel, Jeffrey A. 2012. "The natural resource curse: a survey of diagnoses and some prescriptions."

Gylfason, Thorvaldur. 2011. "Natural resource endowment: A mixed blessing?"

Hamilton, James D. 1989. "A new approach to the economic analysis of nonstationary time series and the business cycle." Econometrica: Journal of the Econometric Society, 357-384.

Hamilton, James D. 2011. Historical oil shocks. No. w16790. National Bureau of Economic Research.

Higgs, Helen and Andrew Worthington. 2008. "Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market." Energy Economics 30, No. 6: 3172-3185.

Hogan, William W. 1983. "Oil stockpiling: help thy neighbor." The Energy Journal 4, No. 3: 49-71.

Kim, Chang-Jin and Charles R. Nelson. 1999. "State-space models with regime switching: classical and Gibbs-sampling approaches with applications." MIT Press Books 1.

Landon, Stuart and Constance Smith. 2015. "Rule-based resource revenue stabilization funds: a welfare comparison." The Energy Journal 36, No. 2: 117-43.

Landon, Stuart and Constance Smith. 2013. "Government revenue stabilization funds: Do they make us better off?." Canadian Public Policy 39, No. 1: 71-99.

Landon, Stuart and Constance E. Smith. 2010. "Energy prices and Alberta government revenue volatility."

Linnemann, Ludger. 2006. "The effect of government spending on private consumption: a puzzle?" Journal of Money, Credit and Banking 38, No. 7: 1715-1735.

Lopez-Murphy, Pablo and Mauricio Villafuerte. 2010. "Fiscal policy in oil producing countries during the recent oil price cycle."

Murphy, Frederic and Fernando S. Oliveira. 2013. "Pricing option contracts on the strategic petroleum reserve." Energy Economics 40: 242-250.

Murphy, F. H., M. Toman, and H. Weiss. 1987. "A nash model of consuming nations' strategic petroleum reserves." Management Science 33, No. 4: 484-499.

Murphy, Frederic H., Michael A. Toman, and Howard J. Weiss. 1986. "An integrated analysis of US oil security policies." The Energy Journal 7, No. 3: 67-82.

Ossowski, Rolando, Steven A. Barnett, James Daniel and Jeffrey M. Davis. 2001. Stabilization and Savings Funds for Nonrenewable Resources. No. 205. International Monetary Fund.

Pereira, Marcelo, Sofia B. Ramos, and José G. Dias. 2016. "The cyclical behaviour of commodities." The European Journal of Finance, 1-22.

Pierru, Axel and Walid Matar. 2014. "The impact of oil price volatility on welfare in the Kingdom of Saudi Arabia: implications for public investment decision-making." The Energy Journal 35, 99-118."

Ross, Michael L. 2015. "What have we learned about the resource curse?" Annual Review of Political Science 18: 239-259.

Saudi Arabia Monetary Authority. 2016. "Annual government revenues and expenditures" Table 2 accessed via: http://www.sama.gov.sa/en-US/EconomicReports/Pages/YearlyStatistics.aspx.

Shenon, Philip. 1995. "A Pacific island nation is stripped of everything." The New York Times.

Sims, Christopher A., and Tao Zha. 2006. "Were there regime switches in US monetary policy?" The American Economic Review 96, No. 1: 54-81.

References

Sturm, Michael, Francois J. Gurtner, and Juan Gonzalez. 2009. "Fiscal Policy Challenges in Oil-Exporting Countries-A Review of Key Issues."

Swensen, David F. 2009. Pioneering portfolio management: An unconventional approach to institutional investment, fully revised and updated. Simon and Schuster.

Teisberg, Thomas J. 1981. "A dynamic programming model of the U.S. strategic petroleum reserve." The Bell Journal of Economics, 526-546.

Villafuerte, Mr Mauricio, Mr Rolando Ossowski, Theo Thomas and Paulo A. Medas. 2008. Managing the Oil Revenue Boom: The Role of Fiscal Institutions. No. 260. International Monetary Fund.

Winston, Wayne L., and Jeffrey B. Goldberg. 2004. Operations research: applications and algorithms. Vol. 3. Belmont: Thomson Brooks/Cole.

Zhang, Yue-Jun and Lu Zhang. 2015. "Interpreting the crude oil price movements: Evidence from the Markov regime switching model." Applied Energy 143: 96-109.

About the Authors

Nader AlKathiri

Nader AlKathiri is a senior research associate at KAPSARC working on the role of stabilization and sovereign wealth funds in oil exporting countries. He holds an MSc in Applied Mathematics and Computational Science and an MBA in Finance.

Tarek Atalla

Tarek is a former research associate evaluating energy productivity investments, economics of energy vulnerability, and the effect of climate on energy consumption patterns.

Frederic Murphy

Frederic is a senior visiting fellow and professor emeritus, Temple University. He has a Ph.D. in Operations Research and a B.A. in Mathematics from Yale University.

Axel Pierru

Axel leads the Energy Systems and Macroeconomics program at KAPSARC. He has published more than 30 papers in peer-reviewed journals. Previously, he worked for 15 years at IFP Energies Nouvelles in France. Axel received his Ph.D. in Economics from Pantheon-Sorbonne University in Paris.

About the Project

Strategies can be developed to mitigate the effects of oil price shocks on the Kingdom's economy. Oil stabilization funds provide short-run protection against oil revenue fluctuations – in the way that Saudi government deposits and reserves at the Saudi Arabian Monetary Authority (SAMA) have historically served as a buffer to decouple government budget from oil revenue fluctuations. Sovereign wealth funds create income for future generations to replace revenue streams from depletable resources – one of the purposes of Saudi Arabia's Public Investment Fund. We have developed a framework for optimizing policies for adding to and withdrawing from stabilization funds, which we applied to Saudi Arabia as a case study. We will pursue this work by creating an integrated framework for linking stabilization funds with sovereign wealth funds. Recent data show the dependence of the non-oil sector's growth to the oil price, whereas the non-oil sector growth is key in achieving the Kingdom's economic diversification objectives. We will study this dependence with a deep dive in data, national accounts and economic theory.

KAPSARC

مركز الملك عبدالله للدراسات والبحوث البترولية

King Abdullah Petroleum Studies and Research Center

www.kapsarc.org