

OLIVEIRA, F.S., ZAHUR, N.B. and WU, F. 2023. Analysis of the optimal policy for managing strategic petroleum reserves under long-term uncertainty: the ASEAN case. *Computers and industrial engineering* [online], 175, article number 108834. Available from: <https://doi.org/10.1016/j.cie.2022.108834>

Analysis of the optimal policy for managing strategic petroleum reserves under long-term uncertainty: the ASEAN case.

OLIVEIRA, F.S., ZAHUR, N.B. and WU, F.

2023



Analysis of the optimal policy for managing strategic petroleum reserves under long-term uncertainty: The ASEAN case

F.S. Oliveira^{a,*}, Nahim B. Zahur^b, Fulan Wu^c

^a Department of Business Analytics, Circular Economy and Entrepreneurship, Bradford School of Management, University of Bradford, Bradford, West Yorkshire, BD7 1DP, UK

^b Department of Economics, Queen's University, Kingston, ON, K7L 3N6, Canada

^c China Academy of West Region Development, Zhejiang University, No. 866, Yuhangtang Road, Hangzhou, Zhejiang 310058, China

ARTICLE INFO

Keywords:

Inventory management
Stochastic programming
Uncertainty
Petroleum

ABSTRACT

We examine the issue of petroleum stockpiling in the Association of Southeast Asian Nations (ASEAN), computing the optimal build-up and draw-down policies under different conditions. We study, in detail, the properties of petroleum prices, oil imports and production, and GDP, analyzing the impact of the planning horizon, discount rate and price elasticity of demand on the optimal policy. We use a finite horizon stochastic program (with varying branching) in which the policymaker minimizes the negative impacts of oil price increases on the GDP and the cost of holding the strategic petroleum reserve. We propose an inter-generational equity rule to compute the level of inventory in the final states of the decision tree. We find that ASEAN countries would benefit significantly from developing a strategic petroleum reserve, with net benefits ranging from US\$25–125 billion. Our suggested target stockpile is consistent with the International Energy Agency's recommendation of holding stocks equal to 90 days of net imports.

1. Introduction

Energy plays a vital role in modern societies, underpinning all areas of economic activity. The economic impact of supply or price disruptions can therefore be significant and wide-ranging (Bohi & Toman, 1993). The oil price shocks of the 1970s, and the resulting macroeconomic disruptions in oil-importing countries, provided the initial impetus for building strategic petroleum reserves. In the aftermath of the 1973 oil crisis, a number of OECD countries established the International Energy Agency (IEA) to coordinate responses to oil market disruptions, with stockpiling of petroleum emerging as one of the major policy tools. These reserves are used to mitigate the risk of shortage (e.g., Rodriguez-Espindola, Alem, & Da Silva, 2020). As such, each IEA member country is required to have a strategic petroleum reserve equal to at least 90 days of its net oil import requirements, where “oil” includes crude oil, natural gas liquids, as well as refined products (International Energy Agency, 2022). Throughout the paper we follow the same convention.

The earliest stockpiling models attempted to assess the value of oil stockpiles by determining the extent to which they could reduce the economic costs of oil supply interruptions, e.g., Hogan (1981), Balas (1981) and Rowen and Weyant (1982). Using dynamic programming models, Teisberg (1981) and Chao and Manne (1983) looked at how to

determine not just the optimal size of the strategic petroleum reserve but also the optimal fill-up and draw-down rates, given information on the duration and frequency of oil supply interruptions. Hogan (1983) explored the interactions between the strategic petroleum reserve (SPR) of a major oil importer and the SPRs of other oil importing countries, paying particular attention to free-riding issues. Murphy, Toman, and Weiss (1986) examined the interaction between public and private stockpiling in dealing with oil disruptions.

More recently, Fan and Zhang (2010) modeled the buildup of petroleum reserves in China and India, studying the interaction between their policies within a game-theoretic framework; Murphy and Oliveira (2010, 2013) developed a Markov game to model the optimal build-up and draw-down of the U.S. strategic petroleum reserve using option contracts; Wu, Fan, Liu, and Wei (2008), Bai, Zhou, Zhou, and Zhang (2012), and Xie, Yan, Zhou, and Huang (2017) studied optimal build-up strategies for China's petroleum reserve, taking into consideration potential market disruptions; Jiao, Han, Wu, Li, and Wei (2014) developed a system dynamics model of China's SPR and use the model to analyze the circumstances under which an SPR release can stabilize the domestic oil price, while Bai, Zhou, Tian, and Meng (2016) analyzed strategic petroleum reserves policy in the presence of supply uncertainty; Bai and Dahl (2018) subsequently evaluated the

* Corresponding author.

E-mail address: f.oliveira@bradford.ac.uk (F.S. Oliveira).

costs and benefits of the US SPR, finding that the net benefits are negative. Finally, Wang, Sun, Li, Chen and Liu (2018) analyzed the optimization of the oil-import portfolio in China, which has been partly used to build the strategic reserves.

As a methodological contribution we develop a recursive decision tree (e.g., Wang & Chuang, 2016) to model uncertainty in multistage stochastic programming (e.g., Azizi, Hu, & Mokari, 2020; Khalilabadi, Zegordi, & Nikbaksh, 2020). We then use ARIMA (e.g., Wang & Chen, 2019) and co-integration, based on historical behavior of oil prices, and other stochastic variables to simulate a wide range of possible future paths for these variables. This allows the optimal stockpiling policy chosen in the dynamic programming model to take into account the level of uncertainty in the oil market. The framework we propose is more flexible than that used in most of the earlier stockpiling literature, where the oil market is typically assumed to oscillate between a discrete number of Markov states. As a result, many of the statistical properties of stochastic variables (such as the random walk behavior commonly observed in oil prices) cannot be incorporated into the policy choice. Our model is solved as a stochastic program, in a finite planning horizon, that minimizes the total social cost of holding the inventory (including both the direct costs associated with holding the stockpile and the indirect cost to GDP of any oil price changes brought about by stockpiling) by deciding, in each period and under different scenarios, how much oil to buy or sell from the petroleum reserve.

We chose GDP as our measure of performance because it has long been identified the link between petroleum price increases and a country's economic output. Hamilton (1983) demonstrated the strong correlation between petroleum price increases and decreases in production in the U.S. economy, for the period 1948–72. Mork (1989) and Mory (1993) reached the same conclusion in subsequent studies. Analyses carried out by Herrera (2009) and Herrera, Lagalo, and Wada (2011) reaffirm the significance of oil price increases for the U.S. economy.

As a policy contribution we therefore analyze the building of a strategic petroleum reserve in the ASEAN (Association of Southeast Asian Nations). This issue has received scant attention thus far, in part because ASEAN has traditionally been rich in natural gas and oil. However, ASEAN has been a net oil importer since 1993, with net oil imports accounting for half of its oil consumption in 2010 (e.g., Nicolas, 2009). We show that ASEAN countries would benefit significantly from risk-pooling to build up a strategic petroleum reserve. The resulting net economic benefits would be at least US\$25–125 billion (depending on the discount factor). At the end of the 25-year planning horizon, The ASEAN's target stockpile is 112 days of net oil imports (with an average size of about 800 MMB) if the social discount factor is 0.99, and 53 days of net oil imports (with an average size of about 380 MMB) if the social discount factor is 0.95. Our suggested target stockpile is consistent with the IEA's requirement that member countries hold stocks equal to 90 days of net imports. This exceeds the current oil stock targets set by ASEAN countries, with the exception of Vietnam which has planned to achieve oil reserves amounting to 90 days of oil imports by 2020. Additionally, our experiments suggest that through cooperative management of the petroleum reserve by adopting a risk pooling strategy, the ASEAN will reduce the maximum stockpile size by about 5000 MMB, a reduction of over 60% of the reserves without risk pooling.

Finally, we are able to identify the pattern for an optimal response to a shock in the petroleum market: the best response to a severe oil price shock involves an immediate large sale of stocks, followed by further sales in future periods, with some stocks reserved to deal with potential future crises. The imposition of the inter-generations rule has only a limited effect on the freedom of the stockpile manager in mitigating the impact of oil price shocks.

The remainder of the paper is organized as follows. Section 2 describes the dynamic programming model of stockpiling management. We discuss the model parameterization, paying particular attention to how we econometrically model the stochastic behavior of oil prices, oil imports, oil production, and GDP. Section 3 presents and analyzes the empirical results from a range of simulations of the stockpiling model developed in this paper. Section 4 concludes the article.

2. Model description

We consider a finite horizon problem represented by a scenario tree. Exogenous stochastic processes determine the oil world price, oil imports and production, and GDP at every node of the tree. The stockpile manager decides how much petroleum to buy or sell at any given node, in order to minimize the sum of the economic cost (measured as the loss of GDP) due to increases in oil prices and the cost of holding the reserve. We consider two alternative stockpiling regimes. In the cooperative case, the ASEAN countries jointly build a regional petroleum reserve that is to be shared by all the countries. In the non-cooperative case, each ASEAN country individually builds its own petroleum reserve, without regard for the effects on other ASEAN countries.

2.1. Simulation of stochastic processes using time-series analysis

The stockpiling strategies depend on several stochastic variables, such as oil prices, oil imports, oil production, and GDP, that are not known to policymakers in advance. We therefore begin by discussing how we simulate the uncertainty in the stochastic variables, before describing the dynamic programming model for making stockpile decisions.

Table 1 describes the variables and data sources. Although the oil price (evaluated in real terms) fluctuates considerably in the short-term, we use the annual average price as we consider long-run stockpile build-up and draw-down policies, e.g., Zhang, Qin, and Chen (2017) and Bai et al. (2012). The Augmented Dickey–Fuller unit root test (see, e.g., Nerlove & Diebold, 1990) indicates that the oil price is a non-stationary series. After taking first differences, the oil price series is stationary. We also find no evidence that oil price changes are serially correlated. As such we model the oil price p_t^o as a random walk, and take the first difference before implementing the regression. Our assumption that oil prices follow a random walk without drift is consistent with the empirical literature on the statistical behavior of historical crude oil prices. For example, Alquist and Kilian (2010) compared a variety of econometric models of the oil price based on not just the current spot price of crude oil, but also survey forecasts of the oil price, current oil futures prices, and the oil futures spread. They found that the model of a random walk without drift provides superior forecast accuracy compared to the alternative models.

The estimated oil price model can thus be represented by Eq. (1), in which $u_t^p \sim N(0, 14.7^2)$. It is estimated based on the historical path of oil prices, which reflect past supply interruptions and, therefore, include severe supply interruptions in the form of large exogenous oil price shocks. Note that there are two identical approaches to modeling the statistical behavior of consumption, production, and net oil imports: treating consumption and production as stochastic variables, with net imports being the difference between consumption and production; and treating net imports and production as stochastic variables, with consumption then equaling the sum of net imports and production. We follow the latter approach. Nonetheless, the results are not affected by the choice of method.

$$p_t^o = p_{t-1}^o + u_t^p \quad (1)$$

We also develop econometric models for real GDP, net oil imports and oil production, as these are all key variables in determining the oil dependency of ASEAN countries. We first describe how we model the behavior of these variables for the ASEAN region as a whole. The Augmented Dickey–Fuller test indicates that i_t , g_t and o_t are non-stationary and integrated of order one (I(1) series). Because of this, we model o_t as a random walk and g_t as a random walk with drift (where the drift term captures the tendency of GDP per capita to grow over time). The models for o_t and g_t are represented by Eqs. (2) and (3) as

Table 1
Variables and Data Sources.

Variable	Data	Source
p_t^0	The real annual crude oil price (measured in 2011 USD per barrel) from 1970 to 2012	BP (2020)
g_t	Real GDP per capita of ASEAN countries (in constant USD) from 1989 to 2010	World Bank (2020)
i_t	Net oil imports per capita of ASEAN countries (in thousand tonnes of oil equivalent or ktoe) from 1989 to 2010	International Energy Agency, 2020
o_t	Oil production per capita of ASEAN countries (in ktoe) from 1989 to 2010	International Energy Agency, 2020
D_t^w	World oil consumption (in million barrels per year)	International Energy Agency, 2020
G_t^w	World real GDP (in million constant USD per year)	World Bank (2020)

follows, where $u_t^o \sim N(0, 0.0084^2)$ and $u_t^g \sim N(0, 39^2)$, and t-statistics are provided in parentheses.

$$o_t = o_{t-1} + u_t^o \quad (2)$$

$$g_t = \underset{(5.45)}{44} + g_{t-1} + u_t^g \quad (3)$$

We next model net oil imports per capita, i_t , as a function of GDP per capita (g_t) (since oil demand increases as the economy grows) and oil production (since oil imports decrease as oil production increases). We regress i_t on g_t and o_t , with the error term modeled as an AR(1) process. The Augmented Dickey–Fuller test shows that the residuals from this equation follow a stationary process, suggesting there is no need to model i_t in differences. The model for i_t is illustrated in Eqs. (4) and (5), where $v_t^i \sim N(0, 0.014^2)$. The coefficients have the expected sign, and are statistically significant (Table A.1 in Appendix A).

$$i_t = \underset{(4.22)}{0.0002}g_t - \underset{(-2.15)}{0.76} o_t + u_t^i \quad (4)$$

$$u_t^i = \underset{(6.26)}{0.76}u_{t-1}^i + v_t^i \quad (5)$$

Eqs. (2)–(5) capture the stochastic behavior of ASEAN’s GDP per capita, net oil imports per capita and oil production per capita, and can be used to generate forecasts and simulations of their behavior in the future. A similar econometric analysis is also carried out for each individual ASEAN country, with details provided in Appendix B.

Finally, it is also necessary to model the world’s oil consumption as it partly determines the impact of a stockpiling intervention on oil prices (as described in Section 2.3). We model world oil consumption D_t^w in conjunction with world real GDP G_t^w , as the two variables are correlated with one another. The two time series are non-stationary and integrated of the same order. We model G_t^w as a random walk with drift, where the drift term represents the long-term evolution of world GDP. The model for G_t^w is illustrated in Eq. (6), where $u_t^G \sim N(0, 430000^2)$.

$$G_t^w = \underset{(11.26)}{720000} + G_{t-1}^w + u_t^G \quad (6)$$

Since long-run oil consumption growth depends on the rate of economic growth, we regress ΔD_t^w on ΔG_t^w , with the error term modeled as an AR(1) process. This is illustrated in Eqs. (7) and (8), where $v_t^D \sim N(0, 400^2)$. The estimates are shown in Table A.2 in Appendix A.

$$\Delta D_t^w = \underset{(5.74)}{0.00073}\Delta G_t^w + u_t^D \quad (7)$$

$$u_t^D = \underset{(5.15)}{0.62}u_{t-1}^D + v_t^D \quad (8)$$

2.2. Procedure used to build scenario trees

Let t and n be, respectively, the index for the time period and the node. Let T represent the number of time periods and N represent the number of nodes. In the scenario tree, when we advance by one time period from time t to $t+1$, each of the nodes at time t (the *parent* nodes) branches out into multiple nodes (which we call *child* nodes), with the latter representing different possible scenarios for the stochastic variables at time $t+1$ (i.e., oil prices, oil imports, oil production and GDP). In building the tree we use a branching factor r_t , i.e., the number of child nodes at time $t+1$ created from any given parent node at time t is r_t . Leaf-nodes are associated with the final period T , and have no descendants. The number of scenarios we consider is therefore $\prod_{t=0}^T r_t$.

The simplest way to construct a scenario tree for the oil stockpiling problem is to have each time period represent a single year and use the same branching factor for every time period. However, because the tree grows exponentially over time, the computational constraints associated with this approach are very large: for instance in a tree with a branching factor of 2 over 30 time periods, the total number of scenarios is 2^{30} or over a billion. This computational burden is too high to be handled with the available computers. Another important reason not to use a brute-force approach (i.e., where the tree grows exponentially over time) is its inefficiency, as only a few states of the world are considered during the initial periods, whereas the number of states considered in the final periods is much greater than required.

For these two reasons, we build a scenario tree in which the branching factor r_t is set higher during the initial periods, so that if $t' < t^*$, $r_{t'} \geq r_{t^*}$. There are three main reasons for this approach. First, the tree grows quickly during the first few periods, which allows a better representation of the stochastic behavior of the variables we are considering, as the impact of the scenarios in the periods closer to the present is more important. Second, the tree grows at a slower rate in the last few periods, as these have a lower impact on the current policy. Even with this approach the number of nodes in the latter stages is very high. Third, we define the time periods t so that they represent shorter time intervals during the initial periods and longer time intervals during the later periods; thus the time interval between periods t and $t+1$ will differ depending on the value of t .

In order to test this approach and build a very large tree, we have tested different branching factors and time horizons. Here we report the results for the approach that gave the highest level of analytical detail, can be computationally solved in reasonable time and yields robust results: our testing showed that an increase in the branching factor and size of the tree would not affect the results significantly.

Fig. 1 illustrates a simplified scenario tree with 2 time periods, 9 nodes and 6 possible scenarios that the decision maker at $n=0$ needs




Fig. 1. Example of a scenario tree with $T = 2$, $N = 9$, $r_1 = 3$, $r_2 = 2$.

to take into account. In our analysis, we considered a total of 4096 scenarios, with the branching factor set at 8 in the first period, 4 in the second period, and 2 in the remaining 7 periods. Note that because we consider a very large number of randomly generated scenarios (over 4000), where the scenarios are drawn based on the statistical behavior of oil prices, oil imports and GDP, the sample size is large enough to be representative of the stochastic environment faced by the policymaker.

Each node in the scenario tree is a state of the world and represents a different realization of the oil price, oil imports, oil production, and GDP. These values are generated stochastically based on the values at the parent nodes and randomly generated exogenous shocks, using the statistical model described in Section 2.1. For instance, suppose at node n_p in time period t , the oil price is $p^o(n_p)$. Since the branching factor is r_t , any one of r_t nodes can be reached from node n_p in period $t + 1$. For any such “child” node n_c , we randomly generate (using the distribution given in Eq. (1)) a normal disturbance $u^o(n_c)$ and calculate the new oil price as $p^o(n_c) = p^o(n_p) + u^o(n_c)$. Similarly, we use Eqs. (2)–(8) to randomly generate regional oil production, regional GDP, net oil imports, world GDP, and world oil consumption in the new child node. We recursively perform this procedure for every node of the tree, thus ending up with N distinct scenarios for oil prices, GDP, net oil imports, and oil production. (Each scenario in this context consists of a 25-year path for each of the random variables.)

It is important to note that the baseline scenarios are based on the statistical models described in Section 2.1. Since these statistical models are necessarily estimated using historical data, they cannot account for potential structural breaks that may occur in the future. For example, if ASEAN economies were to implement stringent climate mitigation policies in the future, this may lower their reliance on petroleum in the future, a possibility which we cannot directly capture in our baseline scenarios. Instead, we carry out sensitivity analysis to check how the optimal behavior of the SPR manager would change if ASEAN economies were to lower their dependence on oil.

The scenario tree is a representation of the stochastic environment within which the policymaker operates. We now turn to the endogenous model describing how the policymaker can optimally make stockpile decisions when faced with this uncertainty .

2.3. Optimizing the petroleum reserve

The manager of the petroleum reserve aims to minimize the sum of the economic cost (measured as the loss of GDP) due to oil price shocks and the cost of holding the reserve, as described by Eqs. (9)–(16), which are defined over the N nodes of the scenario tree. Table 2 summarizes the key variables of the model.

Table 2
Model Variable Definitions.

Variable	Explanation
n	denotes any node of the tree
n_c	child node
n_p	parent node
N	number of nodes
$x(n)$	change in inventory in node n
$x^+(n)$	quantity of oil added to stockpile
$x^-(n)$	quantity of oil sold from stockpile
$y(n)$	inventory in node n at end of period
$k(n)$	stockpile capacity at node n
$M(n)$	minimum target stock level for node n
$p^0(n)$	oil price before stockpile intervention
$p'(n)$	oil price after intervention
$\Delta p(n)$	change in oil price due to intervention
$G(n)$	GDP at node n
$C(n)$	cost incurred in node n
$\Pi(n)$	cost-to-go in node n

For any child node n_c , let n_p denote the corresponding parent node. Let $y(n)$ and $x(n)$ stand for the quantity of oil held in the reserve at node n (in barrels) and the quantity bought for (> 0) or sold from (< 0) the reserve at node n (in barrels), respectively. When we need to distinguish between stockpile purchases and sales, we use the notation $x^+(n)$ to represent the quantity of oil added into the stockpile and $x^-(n)$ to represent the quantity of oil sold from the stockpile. $M(n)$ represents the minimum target level of stocks for the leaf nodes, i.e., nodes corresponding to the final period of the planning horizon. Let $k(n)$ represent the capacity of the oil stockpile at node n (in barrels). The oil price before the stockpile intervention (i.e., the randomly generated oil price) is $p^0(n)$ while the oil price after the intervention is $p'(n)$, with $\Delta p(n)$ denoting the change in oil price due to the intervention.

We start by describing the equation used to compute the level of inventory in the reserve at any given node, Eq. (9). The level of inventory $y(n_c)$ at the end of the period in node n_c equals the final level of inventory in the respective parent node $y(n_p)$ plus the change in the level of inventory, $x(n_c)$.

$$y(n_c) = y(n_p) + x(n_c), \forall n_c \tag{9}$$

We model the building of new stockpile capacity as an endogenous variable, as shown in Eq. (10). This reflects the fact that we need to consider the capital costs of building storage capacity. A similar argument for including the cost of building new capacity is made by Zhu, Liu, and Wang (2012) in the context of China’s SPR: the stockpile capacity at node n_c equals the maximum stockpile capacity in the parent node n_p and the level of inventory in the current node n_c .

$$k(n_c) = \max[k(n_p), y(n_c)], \forall n_c \tag{10}$$

Eq. (11) captures how the oil price changes due to the stockpiling intervention. The oil price before intervention, $p^0(n)$, is randomly simulated for every node using an econometric model for historical oil prices. The change in the oil price due to stockpiling, $\Delta p(n)$, is a function of the oil demand function, which we assume is characterized by a constant elasticity of demand, $\epsilon (< 0)$. The impact of stock changes on the oil price therefore depends on the existing level of prices $p^0(n)$, world oil consumption $D^w(n)$, and the elasticity ϵ . Note that as $\epsilon (< 0)$, it follows from Eq. (11) that an increase (decrease) in the level of the reserve leads to an increase (decrease) in the petroleum price.

$$\begin{aligned} p'(n) &= p^0(n) + \Delta p(n) \\ p'(n) &= p^0(n) - \frac{x(n)p^0(n)}{\epsilon D^w(n)}, \forall n \end{aligned} \tag{11}$$

The cost function (12) encapsulates two main components. The first term represents the loss in GDP due to the change in the petroleum price, which is a cost incurred by the economy as a whole. In (12) α (< 0) represents the oil-price elasticity of the GDP. $\frac{\Delta p(n)}{p(n)}$ is the percentage change in the oil price as a function of shocks in the petroleum price (captured by $p^0(n)$) and considering the change in the level of the reserve (11). $G(n)$ represents the level of GDP at node n in the scenario tree and is randomly simulated using econometric models developed from historical data (as described in Section 2.1). When the price increases due to a disruption in the petroleum market that negatively affects the GDP, as α (< 0). Therefore, a supply disruption increases the price and the cost function.

$$C(n) = -\alpha \frac{\Delta p(n)}{p(n)} G(n) + p'(n)x(n) + \beta(k(n) - k(n_p)) + hy(n) + ux^+(n) + dx^-(n), \forall n \tag{12}$$

An important research question is the relationship between petroleum price changes and GDP. Hamilton’s seminal work (Hamilton, 1983) focuses on price increases because this was the major issue during the period under analysis. Subsequently, Mork (1989), Mory (1993) and Hamilton (1983, 2003) identified an asymmetric effect in the impact of petroleum price increases (decreases) on production. More recently, Herrera et al. (2011, p. 472) concluded that “at the aggregate level, there is no evidence against the hypothesis of symmetric responses to oil price innovations of typical magnitude, consistent with results of Kilian and Vigfusson [Quantitative Economics, 2(3), 419–453 (2011)] for U.S. real GDP.” For these reasons, in Eq. (12) we consider that the impact of petroleum price changes on GDP is symmetrical.

The remaining terms in (12) describe the various costs incurred directly by the stockpile manager. The revenue (cost) associated with selling (buying) petroleum is represented by the second term $\Delta p(n)x(n)$. The cost of building new capacity is given by $\beta(k(n) - k(n_p))$, where β is the cost (in \$ per barrel of capacity) of building one additional unit of capacity and $(k(n) - k(n_p))$ is the capacity added in node n . The holding costs are represented by $hy(n)$, in which h stands for the holding cost per barrel. The costs incurred in filling/refilling the stockpile and drawing down from the stockpile are represented, respectively, by $ux^+(n)$ and $dx^-(n)$, where u is the cost incurred in adding one barrel of oil to the stockpile and d is the cost incurred in withdrawing one barrel of oil from the stockpile.

The decision maker minimizes the present value of the costs incurred during the planning horizon, which is represented in the cost-to-go function (13), where $\Pi(n_p)$ stands for the cost-to-go in any node n_p that is not a leaf and which is the parent of nodes n_c ; δ_{n_p} represents the discount factor and is linked to the discount rate τ_{n_p} by the following equation: $\delta_{n_p} = \frac{1}{1+\tau_{n_p}}$; and $\pi(n_c)$ is the probability of reaching scenario n_c , given that we are departing from node n_p . This cost-to-go function ensures that the decisions in the current period take into account the long-term effects: the smaller the discount factor (equivalently, the larger the discount rate), the smaller the weight placed on the future, and the more short-term focused the decision making will be (with, possibly, lower investment in stocks). On the other hand, the larger the discount factor, the more long-term the government perspective will be (with, possibly, increased inventory size). Typically the discount factors, per year, are above 0.9 and less than 1.

$$\Pi(n_p) = C(n_p) + \delta_{n_p} \sum_{n_c} \Pi(n_c)\pi(n_c), \forall n \text{ not a leaf} \tag{13}$$

Eq. (14) constrains the inventory to be non-negative. We have additionally considered various rules specifying the minimum oil reserves as a function of the net amount of oil (including crude oil, natural gas liquids and refined products) imported per day. In Eq. (15), $M(n)$ stands for the minimum reserve imposed by the rule requiring stocks to equal a given number of days of net imports. Finally, the objective function (16) is the present value of all the costs incurred from the inception of the petroleum reserve at node n_0 to the leaf nodes in the final time

Table 3

Parameter Assumptions for the Base Case with a Single Regional Reserve.

Parameter	Value	Unit	Description
δ_{n_c}	0.99	-	Discount factor
ϵ	-0.067	-	Price elasticity of oil demand
β	5	US\$/barrel-capacity	Cost of building capacity
h	0.227	US\$/barrel	Annual holding costs per barrel
u	0.08	US\$/barrel	Cost of adding oil
d	0.1	US\$/barrel	Cost of withdrawing oil
α	-0.06	-	GDP-oil price elasticity

period T . We minimize the cost-to-go function at the starting node n_0 by choosing the quantities of petroleum to buy for and sell from the reserve in each node of the tree.

$$y(n) \geq 0, \forall n \tag{14}$$

$$y(n) \geq M(n), \forall \text{ leaf } n \tag{15}$$

$$\min_{x(n_1), \dots, x(n_N)} \Pi(n_0) \tag{16}$$

2.4. Case study

We use a scenario tree with a planning horizon of $T = 25$ years, divided into 9 distinct time periods, with the first 5 time periods each representing 1-year intervals and the final 4 time periods each representing 5-year periods. The branching factor r_t is set at 8 for the 1st period, at 4 for the 2nd period, and at 2 for the remaining 7 periods. Thus, we consider a total of $N = 4096$ different scenarios.

We use the dynamic programming model described by Eqs. (9)–(16) to compute optimal stockpiling policies. In analyzing the cooperative regime, we compute stockpiling policies for the ASEAN region acting as a whole. In analyzing the non-cooperative regime, we separately compute stockpiling policies for each country, independently.

The parameters used for the ASEAN are summarized in Table 3. In the base case, we assume a discount factor of 0.99 per year (roughly equivalent to a discount rate of 1%). The price elasticity of oil demand is the long-run price elasticity over the 1990 to 2009 period reported in a study by the IMF (2011) and consistent with the results reported by Hamilton (2009) at about -0.06. The GDP-oil price elasticity, or α , is based on our analysis of ASEAN’s net oil imports bill in relation to its GDP. This impact is asymmetric and is consistent with the result -0.07 reported in Hamilton (2003). Finally, the unit cost of building new capacity β , the unit holding cost h , the marginal cost of filling the stockpile u , and the marginal cost of withdrawing oil d are based on Leiby and Bowman (2000), under the assumption that oil will be stored in salt caverns. As pointed out by Leiby and Bowman (2000), Asia Pacific Energy Research Centre (2002) and Leesombatpiboon (2010), salt caverns for oil storage are a feasible option in Thailand.

The parameters for the individual countries are very similar to the ASEAN parameters, except for the GDP-oil price elasticity, α , which varies substantially across countries, reflecting their varying levels of vulnerability to oil price shocks. Table 4 shows the GDP-oil price elasticity assumptions for the individual ASEAN countries: these estimates are again based on our analysis of each country’s net oil imports and GDP, and are broadly consistent with empirical estimates of the macroeconomic effect of oil price shocks in ASEAN countries (e.g., Abeysinghe, 2001; Cunado & de Gracia, 2005; Chang, Jha, Fernandez, & Jam’an, 2011). As the table shows, 8 of the 10 ASEAN countries are net oil importers whose economies are negatively affected by oil price shocks, and it is these 8 countries that could potentially build their own individual petroleum reserves.

3. Computing a petroleum stockpiling policy for the ASEAN

In this section we examine the key findings from the petroleum stockpiling model for ASEAN comparing the cooperative and non-cooperative regimes under different planning horizons.

Table 4
GDP-oil price elasticity estimates for ASEAN countries.

Country	Value	Net oil importer?
Brunei	0.702%	No
Cambodia	-0.103%	Yes
Indonesia	-0.044%	Yes
Laos	-0.029%	Yes
Malaysia	0.026%	No
Myanmar	-0.005%	Yes
Philippines	-0.069%	Yes
Singapore	-0.154%	Yes
Thailand	-0.107%	Yes
Vietnam	-0.028%	Yes




Fig. 2. Average stockpile size per year as a function of the planning horizon, assuming a discount factor of 0.99/year.

3.1. Regional petroleum reserve under various planning horizons

We first consider the impact of the planning horizon on expected benefit from stockpiling. We initially assume there is no minimum target reserve for the final period of the planning horizon, i.e., $M(n) = 0$. In Fig. 2 we analyze the average stockpile size across all the scenarios: the average stockpile size increases with the planning horizon. The largest average stockpile is about 1280 million barrels (MMB) in the 25-year planning horizon, registered in year 15. There is a pattern of building up the reserve in the early stages and then using it in the final years. Therefore, from Fig. 2 it is evident that without a constraint specifying that the reserve needs to be maintained in the long-term, the stockpile is always emptied in the final period. The growth of the maximum reserve size as a function of the planning horizon is due to increases in population and economic growth, and a decrease in production, with the region becoming increasingly dependent on imported petroleum.

Moreover, this shows that with a finite horizon, it is clearly sub-optimal to have no restriction on the leaf nodes, i.e., we need to incorporate a rule in the model that reflects the long-term value of the oil that remains in reserve in the final leaf nodes of the decision tree. This provides a justification for working with different rules for deciding how much oil to keep in the final stage of the planning horizon.

3.2. A long-term target stockpile size for the regional petroleum reserve

We now consider various rules for the choice of stockpile size in the final period (i.e., the 25th year), discussing a criterion that could be used to guide the choice of the best rule. Such a rule would stipulate the minimum quantity of reserves $M(n)$, as a function of net oil imports, that should be kept in the final stage (refer to Eq. (15)). For example, the IEA uses the 90-day rule, which stipulates that a member state should keep a minimum reserve sufficient to cover 90 days of net imports.

As can be seen in Figs. 3 and 4, the rules lead to an increase in the total reserves kept, not only in the final period but also throughout the planning horizon. The more stringent the rule (i.e., the larger the quantities of petroleum required to be kept in reserve), the greater the level




Fig. 3. Average stockpile per year, using a discount factor of 0.99, as a function of the final stage rule.




Fig. 4. Average stockpile per year, using a discount factor of 0.95, as a function of the final stage rule.

of stockpiling in the earlier periods. For example, with a discount factor of 0.99, the maximum average stockpile per year using the 120-day rule is about 1860 MMB, which represents an approximate increase of 44% in comparison to the maximum average reserve kept when no rule exists. The explanation for this behavior is that the government buys earlier, thus preempting possible price increases during the planning horizon and avoiding last minute stockpiling in the final years to meet the rule requirement (as this would increase prices). A similar pattern holds with different discount factors, although the average size of the stockpile decreases if the discount factor is lower (see Fig. 4).

The total expected cost of stockpiling increases as the rule becomes more stringent. This is not only due to the cost of accumulating oil in order to build a large stockpile, but also because the constraint on the stockpile in the last period reduces the ability of the planner to mitigate the effects of oil market disruptions. Thus, the planner faces a trade-off between the cost of stockpiling and the size of the stockpile in the last period. If the target rule is too stringent, the stockpiling strategies may be entirely dependent on the rule, rather than the optimality conditions. This can lead to the final stockpile size being larger than the stockpile size at any other time purely due to the final period constraint, as illustrated in Fig. 4 with the 120-days rule.

This point is reinforced by Table 5, which shows how the expected cost from stockpiling changes with the stringency of the rule and with the discount factor. When the rule is not very stringent, the cost of building the reserve is lower than the GDP benefits from using the SPR to reduce the impacts of the volatility in oil prices; thus, the expected cost of managing the SPR is negative. When the rule becomes sufficiently stringent, however the cost of building the reserve is higher than the GDP benefits accrued during the current planning horizon, and so the expected cost is positive. The SPR manager thus has to balance the welfare of the future generation (which benefits from a larger stockpile size in the final period) with that of the current generation (which benefits from a lower expected cost).

Rules that result in very limited reserves for future generations are unlikely to be socially optimal. Similarly, highly stringent rules that impose a significant cost on the current generation are not optimal. The question then becomes identifying an adequate rule that is not too stringent but at the same time ensures that in the final stage, an adequate level of reserves is saved for the future generations. We therefore propose the inter-generations rule.

Table 5

Total expected cost (in US\$ million) of managing the petroleum reserve as a function of the final stage rule and of the discount factor.

		Discount Factor	
		0.95	0.99
Final Stage Rule	0 days	-25,414	-124,012
	30 days	-10,581	-85,840
	60 days	3,777	-54,585
	90 days	19,805	-22,475
	120 days	36,063	11,911

Table 6

Rule suggested by inter-generations argument for the 25th year (with different discount factors).

Discount Factor	Rule	Avg. final stock (MMB)
0.9	20-days rule	142
0.95	53-days rule	377
0.99	112-days rule	798

The inter-generations rule: the current generation should not benefit on average from stockpiling and the present value of its investments in stocks should at least equal the discounted value of its GDP benefits. Ideally, the total expected cost from stockpiling should equal zero.

The inter-generations rule ensures that a generation builds a level of stocks such that the potential benefit is equal to the potential cost and leaving the accumulated stock as a safeguard to the next generation. The unconstrained finite horizon optimization will always tend to favor the current generation, at the expense of future generations; the inter-generations rule addresses this tendency and thus serves as insurance for inter-generational equity. Ensuring inter-generational equity in this fashion is necessary because oil, and therefore oil stocks, is an exhaustible resource: at some point in the future it may not be possible to build up an oil stockpile at a reasonable cost and thus the current build-up of oil stocks must take into account the welfare of all future generations.

While we are not aware of any similar criterion for stockpile buildup, the inter-generations rule we propose has similarities to Hartwick's well-known rule for sustainability in the context of resource depletion. Hartwick's rule is also motivated by inter-generational equity concerns, and prescribes reinvesting rents from the extraction of exhaustible resources in capital stocks, thus preventing the current generation from "shortchanging" future generations (Hartwick, 1977).

Table 6 illustrates that for the ASEAN, the rules meeting the inter-generations argument are 20 days (discount factor 0.9), 53 days (discount factor 0.95), and 112 days (discount factor 0.99). These results suggest that the IEA's suggested 90-days rule implies a discount factor between 0.95 and 0.99, and thus seems an adequate rule for ensuring inter-generational equity.

Fig. 5 shows that the maximum average stockpile size reached during the planning horizon is about 280 MMB, 670 MMB and 1790 MMB when the discount factor is 0.9, 0.95 and 0.99, respectively (in each case the maximum is reached in year 15). Both the final stockpile size and the maximum average stockpile size thus increase monotonically with the discount factor. With a larger discount factor the stockpile size increases as the planner places less weight on the current cost of building up the reserve and more weight on the future GDP impact of oil price shocks.

It is important to note that in our model there is a single regional petroleum reserve built by the ASEAN countries, and we do not explicitly account for petroleum reserves built by non-ASEAN countries (such as the US or China). This is similar to most of the prior literature on stockpiling. The interaction between SPRs built by different countries has been studied by Fan and Zhang (2010) who use a game theoretic model to study the SPR choices of China and India. Nonetheless, implicitly, Eq. (11) considers how the interventions from other countries, interact with ASEAN countries to condition the effectiveness of their policies.




Fig. 5. Maximum average stockpile size per year as a function of discounting for a 25-year planning horizon, with the inter-generations rule in place for the final period.

3.3. Comparing the cooperative and non-cooperative regimes

The analysis so far has been confined to the cooperative regime, where the ASEAN countries jointly build a regional petroleum reserve. This has allowed us to focus on identifying the petroleum reserve management strategies that optimize the overall welfare of the region. However, given that any individual ASEAN country could choose to build its own stockpile rather than participate in the regional scheme, it is also important to understand how stockpiling strategies might evolve under a non-cooperative regime. In this section we analyze the case where each country separately builds its own petroleum reserves, assuming a discount rate of 0.99. For ease of comparison with the cooperative case, we assume that each country adopts a rule specifying that the reserves at the end of the planning horizon should equal at least 112 days of net oil imports—in practice, of course, it may well be that different countries adopt different rules for final period reserves.

Figs. 6 and 7 illustrate the build-up and draw-down strategies for the petroleum reserves built by each of the 8 oil-importing countries in ASEAN (Brunei and Malaysia, being net oil exporters for the foreseeable future, do not build any petroleum reserves). The pattern of building up the reserves initially and drawing them down to the minimum level in the later stages is similar to that observed for the regional petroleum reserve. We observe that each country builds a petroleum reserve that is quite large in absolute terms: for instance the smallest reserve, built by Laos, still reaches a maximum average size of 580 MMB in year 15. Countries with larger economies and greater vulnerability to oil price shocks (i.e., larger absolute α) build larger reserves; this is because the effect of oil price shocks on GDP is the highest for these countries. In addition, countries experiencing rapid economic growth, and which are therefore likely to become more dependent on oil imports in the future (such as Vietnam and Indonesia), build larger stockpiles than might be expected from just inspecting their current oil trade balances. This highlights an interesting conclusion from our model: because it takes several years of stockpile build-up before the SPR can be used to mitigate oil market shocks, the expected future level of oil imports matters more for a country's SPR strategy than its current dependence on oil imports.

The above results indicate that left to their own devices, each country could end up building very large stockpile reserves. This is because in the absence of any sharing of reserves or any guarantee that its neighbors will engage in reserve-building, any individual ASEAN country has to rely solely on its own reserves to mitigate the negative macroeconomic impacts of oil price shocks. The lack of coordination in the non-cooperative case leads to duplication of reserve-building and an overall level of reserves within the ASEAN that is much higher than optimally required, as Fig. 8 below illustrates. Thus cooperating to build a joint regional petroleum reserve has important benefits: by coordinating the build-up of reserves and by sharing reserves during oil




Fig. 6. Average stockpile size per year for Indonesia, Philippines, Singapore, Thailand and Vietnam, assuming a discount factor of 0.99/year and a 112-days rule on the leaf nodes.




Fig. 7. Average stockpile size per year for Cambodia, Laos and Myanmar, assuming a discount factor of 0.99/year and a 112-days rule on the leaf nodes.




Fig. 8. Petroleum reserves held within the ASEAN region, with and without cooperation.

market disruptions, ASEAN countries can avoid wasteful duplication of reserves.

It should be noted that the behavioral assumption underlying this analysis is that each ASEAN country manages its reserves without accounting for the reserve-building patterns of its neighbors. This effectively models a situation where policymakers in any individual country adopt a very risk-averse view of the strategic environment: in the absence of cooperation, there is no guarantee that any other country will build reserves or release them at the appropriate time, and therefore the country must build reserves as if no other country in the region is building reserves. In such a setting, the benefit of cooperation is sharing reserves and thereby avoiding too *much* stockpiling. Alternatively, it is also possible that countries are less risk-averse and assume their neighbors are also building reserves, in which case a game-theoretic model of reserve-building will be more appropriate. In such a strategic environment, free-riding considerations will become important, and it is quite possible that there will be too *little* stockpiling in the absence of cooperation. The question of which strategic setting could emerge in practice is an empirical question that will depend greatly on policymaker preferences and priorities. However, given that a strategic petroleum reserve relates to critical national issues of energy security and foreign import dependence, the risk-averse setting (which we have explicitly analyzed) is not implausible. In either setting, a cooperative approach to stockpiling yields evident benefits, which is why our focus in this paper is on the impact of a regional ASEAN reserve where full cooperation is assumed to exist.

3.4. Responses to oil price shocks

Next we explain how the management of the reserve may respond to extreme shocks in the oil price. The analysis focuses on the cooperative case where the ASEAN countries build a regional petroleum reserve; the results obtained for individual country petroleum reserves (under the “non-cooperative regime”) are similar, and omitted for brevity. Out of all the scenarios considered, we identify those in which there is an oil price disruption in order to analyze the optimal response. We consider price shocks of \$100 or more over a 5-year period (this could happen, for instance, if the price rose by \$20 every year for 5 years). Such shocks occur in approximately 5% of all the scenarios.

We analyze the policy response in terms of the sale of stocks following a shock, considering both different discount factors and different rules. With a discount factor of 0.95, we consider both the case with no rule on the leaf nodes and the case with the 53-days rule suggested by the inter-generations argument (Table 7). With a discount factor of 0.99, we similarly consider both the case with no rule on the leaf nodes and the case with the 112-days rule suggested by the inter-generations argument (Table 8).

In general, the response to an oil price shock involves an immediate large stockpile sale, followed by further sales in the following periods. It is not optimal to empty the entire stockpile directly after an oil price shock, since the possibility of another disruption in the next period cannot be ruled out (this is because oil prices follow a random walk and do not exhibit a mean-reverting tendency). Comparing Tables 7 and 8, it is apparent that a higher discount factor results both in larger stock sales (because the accumulated stock is larger) and a smoother pattern of discharge (due to the greater weight placed on future welfare). The magnitude of the response is greater if the shock occurs in the later periods (presumably because of the larger size of the accumulated stockpile in later periods).

The presence of a rule does not have a major effect on the policy response if the shock occurs in years 10 or 15, though a shock in period 20 leads to a somewhat more muted response if the rule is more stringent (due to the need to balance responding to the shock versus maintaining enough stocks in the final period). An important implication is that the inter-generations rule does not adversely affect the ability of the policymaker to respond to severe oil shocks.

3.5. Sensitivity analysis

Next, we conduct sensitivity analysis on some of the major parameters in the stockpiling model. Table 9 summarizes the base case assumptions for key parameters, as well as the alternative assumptions we test in this section. In all the experiments, we assume the discount factor is 0.99 and impose a 112-days rule on the leaf nodes; we note, however, that the results of the sensitivity analysis do not fundamentally change with different discount factors and/or different rules on the leaf nodes.

We first look at the impact of changing the GDP-oil price elasticity, α (Figure 3.8). Our base case assumption for α is based on our analysis of ASEAN’s net oil imports bill in relation to its GDP. However, α may well be higher in absolute terms if the full macroeconomic costs of oil price shocks exceed the direct impact on the oil imports bill (Toman, 1993). Alternatively, it could also be lower if, for instance, ASEAN economies reduce their oil import dependence in the future. This would be the case, for instance, if ASEAN economies were to decarbonize their economy in the future and reduce their reliance on oil and other fossil fuels. Fig. 9 illustrates that the average stockpile size increases with the GDP-oil price elasticity. Intuitively, the greater the GDP-oil price elasticity, the larger the potential loss of GDP from oil price shocks and, therefore, the greater the incentive to stockpile oil.

We next analyze sensitivity to the own-price elasticity of oil demand, ϵ . As shown in Fig. 10, the more elastic the oil market, the

Table 7
Average stockpile sale after an oil price increase of \$100 or more over 5-year period (discount factor of 0.95).

		Discount factor = 0.95					
		No rule on leaf nodes			53-days rule on leaf nodes		
		Year 10	Year 15	Year 20	Year 10	Year 15	Year 20
Stock sale in	Shock occurs in:						
	Year 10	166			169		
	Year 15	26	390		20	378	
	Year 20	17	253	552	-38	211	518
	Year 25	16	187	463	-279	-130	190

Table 8
Average stockpile sale after an oil price increase of \$100 or more over 5-year period (discount factor of 0.99).

		Discount factor = 0.99					
		No rule on leaf nodes			112-days rule on leaf nodes		
		Year 10	Year 15	Year 20	Year 10	Year 15	Year 20
Stock sale in	Shock occurs in:						
	Year 10	275			208		
	Year 15	266	642		194	558	
	Year 20	302	635	926	229	549	814
	Year 25	389	612	939	260	470	824

Table 9
Parameter assumptions for sensitivity analysis.

Parameter	Base Case	Alternative 1	Alternative 2
Price elasticity of demand, ϵ	-0.067	-0.09	-0.05
GDP-oil price elasticity, α	-0.06	-0.09	-0.03




Fig. 9. Average stockpile size as a function of the GDP-oil price elasticity.




Fig. 10. Average stockpile size as a function of the price elasticity of oil demand.

greater the maximum average stockpile size reached during the planning horizon. This is because in a more elastic oil market, stockpile interventions of a given magnitude have a smaller impact on the oil price, and therefore a larger stockpile is needed to guard the economy against oil price shocks. These results are consistent with previous analyses (e.g., Teisberg, 1981). However, it should be noted that this presupposes a constant GDP-oil price elasticity: it may well be that in a more elastic and flexible oil market, oil price shocks will have a smaller effect on GDP, in which case the stockpile required will be smaller.

Finally, we have also carried out sensitivity analysis with respect to the cost of building new capacity, recognizing that not all ASEAN

countries may have the ability to store oil in salt caverns. Doubling the cost of building storage capacity from US\$5/barrel of capacity to US\$10/barrel would cause the maximum average stockpile per year reached to decrease by only 3.3%. Reductions in the cost of capacity similarly have little impact on our findings: halving the cost of building capacity to US\$2.5/barrel would cause the peak stockpile size to increase by 4.3%.

4. Conclusion and policy implications

In this paper, we have used time series analysis, ARIMA and Cointegration to generate a very wide range of future scenarios for oil prices, oil imports, oil production, and population and GDP growth in the ASEAN countries, and built a detailed scenario tree. We then developed a finite horizon stochastic program in which the policymaker chooses the stockpile build-up and draw-down strategies in order to minimize the negative impacts of oil price shocks on GDP and the cost of building the reserve. A key feature of our model is that it allows for inter-generational equity considerations when choosing the optimal strategies for reserve management, making it different from past analyses of SPR policies.

As a major take-way for policy makers, our experiments suggest that ASEAN countries would significantly benefit from developing a regional petroleum reserve. In the absence of a rule for the minimum reserves that need to be held in the final period, SPR strategies will generally be sub-optimal. We compared different decisions rules setting targets for the reserve to be held in the final period, showing that the more stringent the rule, the greater the expected cost of stockpiling to the current generation. We proposed the inter-generations rule for selecting the final period reserve. The rule ensures inter-generational equity by balancing the welfare of the current generation (by mitigating GDP losses from oil disruptions) with the welfare of future generations (by saving a sufficient amount of oil reserves for future usage). Based on this criterion, the ASEAN region's target stockpile, at the end of the 25-year planning horizon, is 112 days of net oil imports (with an average size of about 800 MMB) if the social discount factor is 0.99, and 53 days of net oil imports (with an average size of about 380 MMB) if the social discount factor is 0.95.

Furthermore, from a policy perspective, we analyzed the savings from a risk-pooling strategy for managing the petroleum inventory. We

compared the regional reserve with a non-cooperative regime where each ASEAN country individually builds its own reserve. We found that the lack of coordination leads to a regional stockpile size that is much higher than necessary for optimizing social welfare in the region. From a practitioner’s perspective, this suggests that a key benefit from risk-pooling is avoiding the duplication of reserves that can be shared and allocated optimally by a risk management system. The total savings from the experiments show that the investment in storage capacity can be reduced by over 60%.

As an insight into the energy policy in the region, we also studied the optimal strategies for managing the regional petroleum reserve in case of an oil price shock. We found that the response to an oil price shock involves an immediate large sale of stocks, followed by further sales in future periods. The imposition of the decision rule suggested by the inter-generations argument showed only a limited effect on the freedom of the stockpile manager in mitigating the impact of oil price shocks. We also analyzed how sensitive our findings are to assumptions on price elasticity of oil demand and the macroeconomic impact of oil price increases. As expected, the more elastic the oil market and the greater the macroeconomic impact of oil price increases, the greater the suggested stockpile size.

From a methodological perspective, we have used ARIMA and cointegration to study the co-evolution of petroleum prices, petroleum production and imports in the ASEAN, real GDP per capita, and the world oil consumption and GDP, which we then used to populate a very large scenario tree spanning a planning horizon of up to 25 years. We analyzed how the results are sensitive to the long-term target set for the reserve and proposing the *inter-generations rule* as a way to build sustainable strategic reserves.

In conclusion, from a theory perspective regarding the management of strategic reserves, the new method we have proposed is able to derive new insights on the interaction between the build-up and draw-down policies and long-term objectives of the decision makers, and the inter-generational transfer of welfare. For the decision makers at the ASEAN, our detailed computational model is able to quantify the investment effort required, the potential social benefits of having a reserve, as well as the very sizable savings to be gained from risk-pooling and cooperation among the ASEAN countries. Finally, our analysis shows that the strategic reserve topic is still very relevant in the South East Asia region, and will become even more so as the region is expected to become increasingly dependent on imports, with oil likely to become a very scarce resource in this part of the world.

CRedit authorship contribution statement

F.S. Oliveira: Writing introduction, Conclusion, Literature review, General comments, Proof-reading. **Nahim B. Zahur:** Methodology, Software, Formal analysis, Writing – original draft, Writing – review & Editing. **Fulan Wu:** Methodology, Formal analysis, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgment

The authors acknowledge financial support from the Energy Studies Institute, National University of Singapore.

Table A.1

Regression Analysis. Dependent Variable: Net Imports.

Variable	Coefficient
C	43.44*** (7.96723)
GDP	0.0002*** (4.71E–05)
Oil Production	–0.7565** (0.35215)
AR(1)	0.762*** (0.12183)
$R^2 - Adj.$	0.972
Akaike info criterion	–5.59
Schwarz criterion	–5.39
S.E. of regression	433462
N. Observations	46

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors are in parentheses.

Table A.2

Regression Analysis. Dependent Variable: World Petroleum Consumption.

Variable	Coefficient
C	–126.9 (183.2)
GDP	0.00073*** (1.28E–04)
AR(1)	0.6199*** (0.120471)
$R^2 - Adj.$	0.496
Akaike info criterion	14.87
Schwarz criterion	14.99
S.E. of regression	398.1
N. Observations	45

Note: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$. Standard errors are in parentheses.

Appendix A. Regression tables

See Tables A.1 and A.2.

Appendix B. Econometric models of petroleum imports

The models of net oil imports per capita i_{jt} , GDP g_{jt} and oil production o_{jt} for each of the ASEAN countries (j) are used together with the population forecasts to simulate net oil imports, GDP and oil production. In all the cases, the Augmented Dickey–Fuller tests indicated that i_{jt} , g_{jt} and o_{jt} (for oil producing countries) are I(1) processes.

Brunei. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $g_{jt} = g_{j,t-1} + u_{jt}^g$, $i_{jt} = 0.00021 - \frac{1.09}{(2.38)} g_{jt} + \frac{-46.05}{(46.05)} g_{jt} + \beta_{oj} o_{jt} + u_{jt}^i$, in which $u_{jt}^o \sim N(0, 5.88^2)$, $u_{jt}^g \sim N(0, 1728^2)$ and $u_{jt}^i \sim N(0, 0.78^2)$.

Cambodia. $g_{jt} = 17.08 + \frac{(2.69)}{(2.69)} g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.62}{(3.63)} u_{jt-1}^g + v_{jt}^g$, $i_{jt} = 0.00017 g_{jt} + u_{jt}^i$, $u_{jt}^i = \frac{0.82}{(5.98)} u_{jt-1}^i + v_{jt}^i$, in which $v_{jt}^g \sim N(0, 11.81^2)$, and $v_{jt}^i \sim N(0, 0.0067^2)$.

Indonesia. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $g_{jt} = 20.69 + \frac{(3.5)}{(3.5)} g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.39}{(2.87)} u_{jt-1}^g + v_{jt}^g$, $i_{jt} = 0.00024 g_{jt} - \frac{0.96}{(5.15)} o_{jt} + u_{jt}^i$, $u_{jt}^i = \frac{0.75}{(5.42)} u_{jt-1}^i + v_{jt}^i$, where $u_{jt}^o \sim N(0, 0.035^2)$, $v_{jt}^g \sim N(0, 25.36^2)$ and $v_{jt}^i \sim N(0, 0.016^2)$.

Laos. $g_{jt} = 19.18 + \frac{(2.27)}{(2.27)} g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.79}{(5.39)} u_{jt-1}^g + v_{jt}^g$, $i_{jt} = i_{jt-1} + u_{jt}^i$, where $v_{jt}^g \sim N(0, 8.05^2)$ and $u_{jt}^i \sim N(0, 0.0007^2)$.

Malaysia. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $u_{jt}^o = \frac{0.31}{(1.97)} u_{jt-1}^o + v_{jt}^o$, $g_{jt} = 89.26 + \frac{(5.65)}{(5.65)} g_{j,t-1} + u_{jt}^g$, $i_{jt} = 0.00015 g_{jt} - \frac{0.86}{(7.12)} o_{jt} + u_{jt}^i$, $u_{jt}^i = \frac{0.057}{(4.67)} u_{jt-1}^i + v_{jt}^i$, in which $v_{jt}^o \sim N(0, 0.095^2)$, $u_{jt}^g \sim N(0, 112.79^2)$ and $v_{jt}^i \sim N(0, 0.057^2)$.

Myanmar. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $u_{jt}^o = \frac{0.36}{(2.37)}u_{j,t-1}^o + v_{jt}^o$, $g_{jt} = g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.98}{(15.53)}u_{j,t-1}^g + v_{jt}^g$, $i_{jt} = 0.015 - \frac{0.32}{(2.1)}o_{jt} + \frac{u_{jt}^i}{(-1.97)}$, $u_{jt}^i = \frac{0.9}{(12.1)}u_{j,t-1}^i + v_{jt}^i$, in which $v_{jt}^o \sim N(0, 0.003^2)$, $v_{jt}^g \sim N(0, 6.58^2)$ and $v_{jt}^i \sim N(0, 0.0035^2)$.

Philippines. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $g_{jt} = g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.49}{(3.93)}u_{j,t-1}^g + v_{jt}^g$, $i_{jt} = \frac{0.000018}{(2.36)}g_{jt} + \frac{u_{jt}^i}{(18.49)}$, $u_{jt}^i = \frac{0.99}{(18.49)}u_{j,t-1}^i + v_{jt}^i$, where $u_{jt}^o \sim N(0, 0.005^2)$, $v_{jt}^g \sim N(0, 27.69^2)$ and $v_{jt}^i \sim N(0, 0.016^2)$.

Singapore. $g_{jt} = \frac{6.13}{(5.47)} + g_{j,t-1} + u_{jt}^g$, $i_{jt} = \frac{2.6}{(3.29)} + \frac{0.00032}{(8.2)}g_{jt} + u_{jt}^i$, $u_{jt}^i = \frac{0.59}{(4.35)}u_{j,t-1}^i + v_{jt}^i$, where $u_{jt}^g \sim N(0, 801^2)$ and $v_{jt}^i \sim N(0, 0.86^2)$.

Thailand. $o_{jt} = \frac{0.065}{(4.98)}o_{j,t-1} + u_{jt}^o$, $g_{jt} = 47.21 + \frac{g_{j,t-1}}{(3.6)} + u_{jt}^g$, $u_{jt}^g = \frac{0.32}{(2.31)}u_{j,t-1}^g + v_{jt}^g$, $i_{jt} = \frac{0.00035}{(6.5)}g_{jt} - \frac{1.6}{(-3.81)}o_{jt} + \frac{u_{jt}^i}{(9.51)}$, $u_{jt}^i = \frac{0.86}{(9.51)}u_{j,t-1}^i + v_{jt}^i$, where $u_{jt}^o \sim N(0, 0.0081^2)$, $v_{jt}^g \sim N(0, 63.26^2)$ and $v_{jt}^i \sim N(0, 0.024^2)$.

Vietnam. $o_{jt} = o_{j,t-1} + u_{jt}^o$, $u_{jt}^o = \frac{0.36}{(2.24)}u_{j,t-1}^o + v_{jt}^o$, $g_{jt} = \frac{30.77}{(1.91)} + g_{j,t-1} + u_{jt}^g$, $u_{jt}^g = \frac{0.92}{(11.89)}u_{j,t-1}^g + v_{jt}^g$, $i_{jt} = -\frac{0.03}{(-9.47)} + \frac{0.0003}{(22.93)}g_{jt} - \frac{0.96}{(-37.39)}o_{jt} + u_{jt}^i$, where $v_{jt}^o \sim N(0, 0.013^2)$, $v_{jt}^g \sim N(0, 4.86^2)$ and $v_{jt}^i \sim N(0, 0.0067^2)$.

References

- Abeysinghe, T. (2001). Estimation of direct and indirect impact of oil price on growth. *Economic Letters*, 73, 147–153.
- Alquist, R., & Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4), 539–573.
- Asia Pacific Energy Research Centre (2002). Energy security initiative: Emergency oil stocks as an option to respond to oil supply disruptions, 2002.
- Azizi, V., Hu, G., & Mokari, M. (2020). A two-stage stochastic programming model for multi-period reverse logistics network design with lot-sizing. *Computers & Industrial Engineering*, 143, Online.
- Bai, Y., & Dahl, C. (2018). Evaluating the management of US strategic petroleum reserve during oil disruptions. *Energy Policy*, 117, 25–38.
- Bai, Y., Zhou, P., Tian, L., & Meng, F. (2016). Desirable strategic petroleum reserves policies in response to supply uncertainty: a stochastic analysis. *Applied Energy*, 162, 1523–1529.
- Bai, Y., Zhou, D. Q., Zhou, P., & Zhang, L. B. (2012). Optimal path for China's strategic petroleum reserve: A dynamic programming analysis. *Energy Economics*, 34, 1058–1063.
- Balas, E. (1981). The strategic petroleum reserve: how large should it be? In B. A. Bayraktar, & et al. (Eds.), *Energy policy planning*. New York: Plenum Press.
- Bohi, D. R., & Toman, M. A. (1993). Energy security: externalities and policies. *Energy Policy*, 21(11), 1093–1109.
- BP (2020). Statistical review of world energy 2020.
- Chang, Y., Jha, K., Fernandez, K. M., & Jam'an, N. F. (2011). Oil price fluctuations and macroeconomic performances in asian and oceanic economies. In *Proceedings of the 30th USAEE/IAEE North American conference*.
- Chao, H. P., & Manne, A. S. (1983). Oil stockpiles and import reductions: A dynamic programming approach. *Operations Research*, 31(4), 632–651.
- Cunado, J., & de Gracia, F. P. (2005). Oil prices, economic activity and inflation: evidence for some Asian countries. *The Quarterly Review of Economics and Finance*, 45, 65–83.
- Fan, Y., & Zhang, X.-B. (2010). Modelling the strategic petroleum reserves of China and India by a stochastic dynamic game. *Journal of Policy Modeling*, 32, 505–519.
- Hamilton, J. D. (1983). Oil and the macroeconomy since world war II. *Journal of Political Economy*, 91, 28–248.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Economics*, 113(2), 363–398.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007-08. *Brookings Papers on Economic Activity*, 40(1), 215–283.
- Hartwick, J. M. (1977). Intergenerational equity and investing rents from exhaustible resources. *American Economic Review*, 67, 972–974.
- Herrera, A. M. (2009). Oil price shocks, systematic monetary policy, and the great moderation. *Macroeconomic Dynamics*, 13, 107–137.

- Herrera, A. M., Lagalo, L. G., & Wada, T. (2011). Oil price shocks and industrial production: Is the relationship linear? *Macroeconomic Dynamics*, 15, 472–497.
- Hogan, W. (1981). Import management and oil emergencies. In D. Deese, & J. Nye (Eds.), *Energy and security*. Cambridge, Mass: Ballinger Publishing Company.
- Hogan, W. (1983). Oil stockpiling: Help thy neighbor. *Energy Journal*, 4(3), 49–71.
- International Energy Agency (2022). Oil security. from <https://www.iea.org/areas-of-work/ensuring-energy-security/oil-security>. (Accessed 3 June 2022).
- International Monetary Fund (2011). World economic outlook 2011: tensions from the two-speed recovery - unemployment, commodities, and capital flows. Washington D.C.
- Jiao, J. L., Han, K. Y., Wu, G., Li, L. L., & Wei, Y. M. (2014). The effect of an SPR on the oil price in China: a system dynamics approach. *Applied Energy*, 133, 363–373.
- Khalilabadi, S. M. G., Zegordi, S. H., & Nikbakhsh, E. (2020). A multi-stage stochastic programming approach for supply chain risk mitigation via product substitution. *Computers & Industrial Engineering*, 149, Online.
- Leesombatpiboon, Poonpat (2010). *Thailand's energy security: strategic petroleum reserve and its economic impacts* (Dissertation), George Washington University.
- Leiby, Paul N., & Bowman, David (2000). *The value of expanding Asian pacific strategic oil stocks*. Oak Ridge National Laboratory, ORNL-1999/39.
- Mork, K. A. (1989). Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy*, 97, 740–744.
- Mory, J. F. (1993). Oil price and economic activity: Is the relationship symmetric? *Energy Journal*, 14(4), 151–161.
- Murphy, F. H., & Oliveira, F. S. (2010). Developing a market-based approach to managing the US strategic petroleum reserve. *European Journal of Operational Research*, 206(2), 488–495.
- Murphy, F. H., & Oliveira, F. S. (2013). Pricing option contracts on the strategic petroleum reserve. *Energy Economics*, 40, 242–250.
- Murphy, F. H., Toman, M. A., & Weiss, H. (1986). An integrated analysis of U.S. oil security policies. *Energy Journal*, 7(3), 67–82.
- Nerlove, M., & Diebold, F. (1990). Unit roots in economic time series: a selective survey. In T. Bewley (Ed.), *Advances in econometrics. Vol. 8*. New York: JAI Press, 1990.
- Nicolas, F. (2009). *ASEAN energy cooperation - an increasingly daunting challenge*. Paper from Institut Francais des Relations Internationales (IFRI, available at IFRI: www.ifri.org/downloads/ nicolas.pdf.
- Rodriguez-Espindola, O., Alem, D., & Da Silva, L. P. (2020). A shortage risk mitigation model for multi-agency coordination in logistics planning. *Computers & Industrial Engineering*, 148, Online.
- Rowen, H. S., & Weyant, J. P. (1982). Reducing the economic impacts of oil supply interruptions: An international perspective. *Energy Journal*, 3(1), 1–34.
- Teisberg, T. J. (1981). A dynamic programming model of the U.S. strategic petroleum reserve. *The Bell Journal of Economics*, 12(2), 526–546.
- Toman, M. A. (1993). The economics of energy security: theory, evidence, policy. In A. V. Kneese, & J. L. Sweeney (Eds.), *Handbook of natural resource and energy economics, III*. Elsevier Science Publishers.
- Wang, C.-H., & Chen, J.-Y. (2019). Demand forecasting and financial estimation considering the interactive dynamics of semiconductor supply-chain companies. *Computers & Industrial Engineering*, 138, Online.
- Wang, C.-H., & Chuang, J.-J. (2016). Integrating decision tree with back propagation network to conduct business diagnosis and performance simulation for solar companies. *Decision Support Systems*, 81, 12–19.
- Wang, J., Sun, X., Li, J., Chen, J., & Liu, C. (2018). Has China's oil-import portfolio been optimized from 2005 to 2014? A perspective of cost-risk tradeoff. *Computers & Industrial Engineering*, 126, 451–464.
- World Bank (2020). World development indicators. Available from <http://data.worldbank.org/data-catalog/world-development-indicators>.
- Wu, G., Fan, Y., Liu, L. C., & Wei, Y. M. (2008). An empirical analysis of the dynamic programming model of stockpile acquisition strategies for China's strategic petroleum reserve. *Energy Policy*, 36(4), 1470–1478.
- Xie, N., Yan, Z., Zhou, Y., & Huang, W. (2017). China's optimal stockpiling policies in the context of new oil price trend. *Energy Policy*, 105, 332–340.
- Zhang, X. B., Qin, P., & Chen, X. L. (2017). Strategic oil stockpiling for energy security: The case of China and India. *Energy Economics*, 61, 253–260.
- Zhu, Z., Liu, L., & Wang, J. (2012). Optimization of China's strategic petroleum reserve policy: A Markovian decision approach. *Computers & Industrial Engineering*, 63(3), 626–633.