
Optimal Strategic Petroleum Reserve Policies: A Steady State Analysis

Author(s): Shmuel S. Oren and Shao Hong Wan

Reviewed work(s):

Source: *Management Science*, Vol. 32, No. 1 (Jan., 1986), pp. 14-29

Published by: [INFORMS](#)

Stable URL: <http://www.jstor.org/stable/2631396>

Accessed: 13/09/2012 14:42

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at
<http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Management Science*.

OPTIMAL STRATEGIC PETROLEUM RESERVE POLICIES: A STEADY STATE ANALYSIS*

SHMUEL S. OREN AND SHAO HONG WAN

*Department of Industrial Engineering and Operations Research, University of California,
Berkeley, California 94720*

*Department of Engineering-Economic Systems, Stanford University,
Stanford, California 94305*

A simple model is presented which allows us to determine the optimal size, fillup, and drawdown rates for a Strategic Petroleum Reserve (SPR) under a variety of supply and demand conditions. The optimal policy variables are determined by minimizing an analytic expression which we derive for the expected insecurity cost rate to the U.S. due to uncertainty in supply of imported oil.

The oil market is modeled in terms of an elastic demand curve and two levels of supply which alternate according to a stationary, continuous time Markov process. The SPR policy is characterized by a fixed fillup rate up to the reserve's capacity, when supply is at its normal level, and a fixed drawdown rate during shortages. The insecurity cost rate being minimized includes consumer welfare loss due to price hikes, reserve holding cost and capital appreciation of the reserve.

Base case results and sensitivity analysis are presented and compared to results obtained by previous approaches. These comparisons suggest that the proposed model can reasonably approximate the more computationally demanding stochastic dynamic programming formulation. The main advantage of the new approach is that it permits extensive sensitivity analysis which is important given the quality of the data.

(INVENTORY/PRODUCTION—STOCHASTIC MODELS, APPLICATIONS; DYNAMIC PROGRAMMING—APPROXIMATION)

1. Introduction

The oil embargo of 1973 and the shut down of the Iranian oil supply during the Iranian revolution in 1979 triggered a wave of studies concerned with energy security. Oil stockpiling policies have been the primary focus of these studies and numerous analytic models have been developed to analyze such policies. These models vary in their complexity, analytic approach, and aspects of the problem emphasized. Early analyses in this area focused on determining the optimal size of a strategic oil reserve using simple static two period models (e.g. FEA 1974, CBO 1980, Hogan 1981, Rowen and Weyant 1982). In these studies the fillup and drawdown policy is predetermined (usually fillup in first period and use up the stockpile in the second period). The optimal stockpile is then determined by evaluating the cost and benefits of stockpiling different amounts of oil with and without a supply interruption. The optimal SPR size is calculated so as to maximize the expected net benefit taking the probability of an interruption into account.

A different type of static model introduced by Balas 1981 focuses on the strategic role of oil stockpiling in the competitive interaction between oil users and oil suppliers. This analysis views the SPR primarily as a deterrent to oil interruption and uses game theory to calculate the optimal SPR size that will provide such a deterrence function. A major shortcoming of both types of static analyses is that they fail to address the stockpile's fillup and drawdown rates. Such policy instruments can be addressed in a

* Accepted by Warren E. Walker; received October 17, 1983. This paper has been with the authors 3 months for 2 revisions.

dynamic framework which takes into consideration the uncertainty in the duration of supply interruptions and the frequency of such interruptions. Teisberg (1981) and Chao and Manne (1982) developed stochastic dynamic programming models that allow explicit consideration of such uncertainty. These models have been used to determine both the size of the SPR and the optimal fillup and drawdown rate contingent upon the supply conditions, time, and available inventory. These models also allow consideration of additional emergency policy instruments such as quotas and tariffs. In a more recent study Hogan (1982) extended Teisberg's model to include gaming aspects which account for possible interactions between the U.S. SPR policy and the SPR policies of other oil importing nations.

While the stochastic dynamic programming models are quite comprehensive, their usefulness is limited by their computational requirements. The optimal policy generated by such models for any given set of assumptions is given as an array of actions contingent upon the state of the world oil market, time, and inventory level. Such results are useful when the objective is to develop contingency plans for operational purposes. However, the large action space and vast number of possible interruption scenarios on which these actions are contingent, tend to confound the broader policy questions concerning the need for an SPR and its impact on the economy and world oil market. When addressing such broad issues, or when SPR policies are examined within a more general energy modeling framework, operational details are of little value. More important is the capability to perform extensive sensitivity analysis which will examine the results under wide variation in the key parameters. Such analysis is essential in this problem given the limited data available. The stochastic dynamic programming approach does not meet this need, since any sensitivity analysis must be represented with reference to specific interruption scenarios. A more suitable approach for the above purposes is to take a macroscopic view of the operational details by limiting the action space, while characterizing the market in terms of a few key parameters that can be easily varied over wide ranges. The need for such an approach has been also recognized in a recent paper by Chao (1983) who presents an elegant theoretical analysis examining the general question of optimal inventory policy in the face of uncertain supply.

The object of this paper is to introduce a new model for analyzing oil stockpiling policies which is comparable in its computational simplicity with the simple two-period models mentioned earlier, and yet captures the important features of the stochastic dynamic programming formulation. In particular, our model captures the uncertainty in the duration of interruptions and time between interruptions. It also takes into consideration the magnitude of the interruption, demand elasticity, fillup rate, and drawdown rate. The objective of our analysis is to determine a SPR policy which will minimize the expected time averaged insecurity cost to the U.S. economy due to uncertainty in the supply of imported oil. This cost consists of three components: the net losses in consumer welfare due to price changes induced by supply interruptions, fillup, and drawdown of the SPR; reserve holding cost; and the net capital appreciation of the stock.

In our model, which may be regarded as a "stripped" version of Teisberg's (1981) formulation, we make two major simplifying assumptions. We assume a stationary environment and restrict the fillup and drawdown rates to be constant (i.e., independent of the reserve level unless, of course, the SPR is empty or full). The stationary assumption means that both supply and demand are independent of time. This neglects the effect of economic growth on demand and the potential increases in the normal price of oil due to its exhaustible nature. While the demand for oil is represented by a general deterministic demand function, supply is assumed to be inelastic and uncertain, with two possible levels corresponding to the "normal" and

disrupted states of the oil market. The transition between states is modeled as a continuous time stationary Markov process. Due to the above simplifications, the expected insecurity cost approaches a stationary growth rate for which we obtain an analytic expression in terms of the SPR policy parameters. The long-run expected time average of the insecurity cost equals that rate. Thus, the long-run optimal policy parameters can be determined by means of static numerical optimization techniques. The simplicity of this computational procedure allows an extensive parametric analysis which provides a comprehensive picture of the solution and its sensitivities with respect to the various assumptions. This capability is very important given the "poor" quality of the available data and it is the main advantage of the new approach over the more elaborate stochastic dynamic programming formulation.

In §2 we describe in more detail the stochastic supply model and obtain a steady state probability distribution for the various states of supply as a function of the stockpiling policy variables. §3 describes the demand model and formulation of the objective function. It also discusses the implications of assuming demand and price stationarity. §4 presents parametric analyses of the optimal stockpiling policy and compares a base case solution to earlier results cited in the literature. We conclude in §5 with some remarks.

2. Supply Model

We assume two possible states of world oil supply: a "normal state" and an "interruption state." The supply in the normal state can be characterized by a supply function $S_N(p)$ where p is the prevailing world oil price and S_N is the annual rate of traded oil on the world oil market. A disruption in world oil supply is described by a fixed downward shift in the supply curve, i.e., $S_I(p) = S_N(p) - Y$ where Y is the decrease in the supply rate. We assign fixed (inelastic) supply levels in each state in our analysis. However, we pursue the formulation with the more general supply characterization. As in the more detailed stochastic dynamic programming formulation by Teisberg (1981), Chao and Manne (1982), and Hogan (1982), we assume that the transitions between supply states can be represented by a stationary Markov process. In a discrete time formulation this would imply a geometric distribution of both the time between interruptions and the duration of an interruption. Here, however, we use a continuous time formulation and hence characterize these transitions by a continuous time Markov process which implies that both the time between interruptions and duration of the interruption are exponentially distributed. (Note that the exponential distribution is just the limiting case of the geometric, as the time increments approach zero.)

Admittedly, the lack of memory assumed by such a model is a rather strong assumption and quite unrealistic, hence one should view such a model as a crude first order approximation of the recurrence aspect in supply interruptions. Conceptually it is possible to add memory to the process by increasing the state space. However, such an approach is limited by data availability and our desire to keep the analysis simple. In fact, the crude nature of the Markovian assumption provides further justification for keeping the operational plan for the SPR and other aspects of the model simple. Optimizing an elaborate operational plan over a long time horizon seems incompatible with the crude nature of the supply model employed in the various dynamic programming formulations.

It is appropriate to mention at this point that while Teisberg's (1981) supply model is stationary, in the sense that the transition probabilities between supply states are independent of time, it is not purely Markovian, since it switches between two transition matrices depending on the reserve level. This feature of Teisberg's model is designed to capture the deterrence effect of an SPR, but we exclude it for two reasons.

First, we have formulated a simple model in order to be able to derive analytic results. Second, we feel that while the above feature of Teisberg's model may capture, qualitatively, the deterrence effect of an SPR, determining the reserve level at which the switchover occurs and the transition probabilities for each case is a highly speculative task. Hence, without extensive sensitivity analysis, which is practically impossible in a dynamic programming framework, such a complication could only provide some qualitative inferences. We believe, therefore, that the deterrence aspect should be analyzed separately in a game theoretic framework.

Let $N(t)$ denote the cumulative probability for an uninterrupted supply period of length t or less and $I(t)$ the cumulative probability for an interruption of duration t or less. Then

$$N(t) = 1 - \exp(-\lambda t) \quad \text{and} \quad I(t) = 1 - \exp(-\mu t),$$

where $1/\lambda$ and $1/\mu$ are the respective mean durations of normal and disrupted supply. It is well known from renewal theory (see Cinlar 1975) that the steady state probability of being in the normal state is then $s = \mu/(\lambda + \mu)$ and the probability of being in the disrupted supply state is $1 - s = \lambda/(\lambda + \mu)$.

A stockpiling policy is aimed at mitigating the adverse effects of disruption by accumulating a reserve during normal supply and releasing it when supply is disrupted. We will assume that the reserve has a maximum capacity C . During normal supply it is filled at a fixed rate U until it reaches its capacity limit. Then when disruption occurs, the stock is released at a fixed rate V until it is depleted or the disruption ends. Figure 1 illustrates a typical pattern of the stockpile level as a function of time. The variables C , U and V are the stockpiling policy variables which we need to determine. The dynamic process describing the combination of supply state and stockpile level can be characterized by two renewal equations defining the probability distribution over the stockpile level in each of the two supply states. This process has been analyzed in a recent article by Meyer, Rothkopf and Smith (1979). In particular they calculate the average stockpile level \bar{L} which is

$$\bar{L} = C \left\{ 1 - \frac{\rho - (1 - s + \rho s)(e^w - 1)/w}{\rho - e^w} \right\} \quad \text{where}$$

$$\rho = \lambda V / \mu U, \quad s = \mu / (\lambda + \mu) \quad \text{and} \quad w = C(\mu / V - \lambda / U).$$

(The above expression for \bar{L} contains a typographic correction to the original expression provided by Meyer, Rothkopf and Smith 1979.)

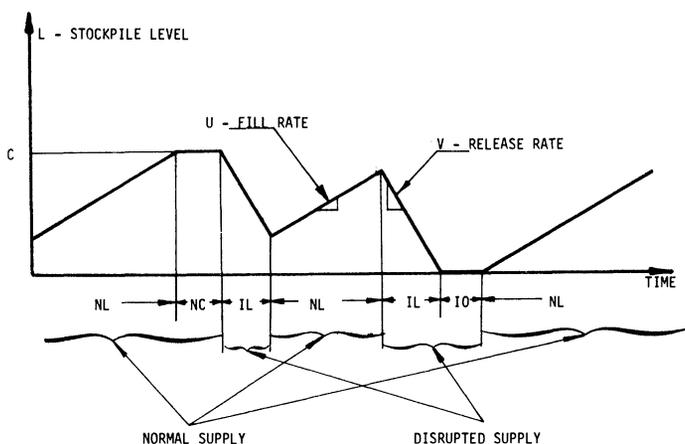


FIGURE 1. Supply States and the Stockpile Level as Function of Time.

TABLE I
Definition of Mitigated Supply States and Corresponding Supply Curves

State	Reserve Level	Supply State	Supply Curve
NC	full ($L = C$)	normal	$S_N(p)$
NL	filling ($L < C$)	normal	$S_N(p) - U$
IO	empty ($L = 0$)	interrupted	$S_N(p) - Y$
IL	releasing ($L > 0$)	interrupted	$S_N(p) - Y + V$

The stockpiling policy described above results in four possible states of mitigated supply which are illustrated in Figure 1. Table I provides a formal definition of each of these states and the corresponding supply curves.

Following Meyer, Rothkopf and Smith (1979), the stationary probabilities for the states NL and IO are given by:

$$P_{NL} = s \{1 - (\rho - 1)e^w / (\rho - e^w)\} \quad \text{and} \quad P_{IL} = (1 - s) \{(1 - e^w) / (\rho - e^w)\}.$$

Clearly then

$$P_{NC} = s - P_{NL}, \quad \text{and} \quad P_{IO} = (1 - s) - P_{IL}.$$

One can also easily verify that $P_{NL}U = P_{IL}V$. The latter is the steady state material balance equation for the reserve (average inflow = average outflow).

3. Objective Function and Demand Model

The main objective of our analysis is to determine the values of an SPR policy parameters that minimize the expected cost incurred by the U.S. economy as a whole due to oil supply insecurity. This cost consists of three components: the net welfare loss to U.S. consumers due to mitigated supply interruptions and SPR fillup, plus the holding cost for the reserve, less the capital gains due to appreciation of the reserve.

In our analysis we ignore issues related to stock financing, private stocks behavior and the distribution of benefits. All these are distributional issues associated with transfer payments within the economy which cancel out when we consider the U.S. economy as a whole. Once the optimal SPR policy is determined, private stocks behavior can be taken into consideration in determining the government owned reserve size and its reserve management policy. The only assumption we need to make at this point is that the private stocks do not exceed the optimal reserve level for the economy. This is a reasonable assumption since risk neutral agents will stock only up to the point where holding cost exactly equals the capital appreciation of the reserve. The government's objective, on the other hand, also includes the consumer welfare. We expect, therefore, that the optimum policy for the economy as a whole will be to stock more than the break-even point, which will, in turn, discourage private stock holdings by risk neutral agents. We verify this assertion in our numerical computations.

Since we consider a stationary environment, our objective function is the expected steady state rate of benefits resulting from a long-term SPR policy. To obtain an analytic representation of this objective function we first need to introduce a demand model and a measure of consumer gains due to shortage mitigation.

A common and undisputed assumption throughout the literature dealing with oil stockpiling is that there is a single world oil market. Hence, any supply interruption and preventive measures will affect the entire world oil market and can be evaluated in terms of their impact on the world oil price. Consequently, some of the benefits resulting from a price decrease through the release of U.S. oil reserves during an interruption will be enjoyed by the other oil importers as well. By the same token,

however, the burden resulting from the price increase due to additional demand during fillup will be shared by the rest of the world. This phenomenon is known as the “free rider” problem which may suggest that a nation could rely on the stockpiles of other nations to alleviate the hardship of a supply interruption. Hogan (1982) has studied the incentives for any nation to maintain its own SPR in spite of the opportunity for a “free ride.” In our model we do not include the game theoretic aspects accounted for by Hogan (1982), but we account for the fact that a U.S. SPR policy will directly affect U.S. consumers only through its effect on the world oil price. To consider the effect of U.S. import level on world oil prices we assume a demand function $D(p)$ describing the annual rate at which oil is demanded on the world oil market at price p . We further assume that U.S. imports constitute a fixed fraction of that demand, i.e. $D_{US}(p) = \gamma D(p)$.

The consumer welfare implication of a price change in world oil is commonly measured in terms of the integral of the demand function over the price change. This quantity is referred to in the economic literature as the change in “consumer surplus” and it approximates the change in consumption utility of utility maximizing consumers (see, e.g., Varian 1984). This measure is based on the economic principle that in an efficient market the price p given by the inverse demand curve $p = D^{-1}(q)$ represents the marginal value to the consumers of the q th unit of consumption. Hence, if consumers can buy all the q units at price p they will break even on the marginal unit q while on all other units they retain a benefit (consumer surplus) which is the difference between the respective economic values, as given by the demand curve, and the price paid p . The total consumer surplus at price p is therefore given by the area under the demand curve to the right of p . Similarly, the difference in consumer surplus due to a price change (assuming instantaneous market response) is given by the integral under the demand curve from the old up to the new price (see illustration in Figure 2).

Assuming that the world oil prices in each supply state is the market clearing price, we can calculate the relative consumer surplus associated with each supply state, as shown in Figure 2. In the absence of a stockpiling policy the oil price under normal supply is p^* and during an interruption of magnitude Y is p_{IO} . The consumer surplus loss during an interruption is then given by the area between p^* and p_{IO} under the curve $D(p)$. The stockpiling policy will cause an increase in price from p^* to p_{NL} during fillup and a reduction in price from p_{IO} to p_{IL} during release. This will result in a consumer surplus loss represented by the cross hatched area between p^* and p_{NL} and a gain represented by the cross hatched area between p_{IO} and p_{IL} . In order to compute the expected benefit of a stockpiling policy we use the normal price level p^* as a

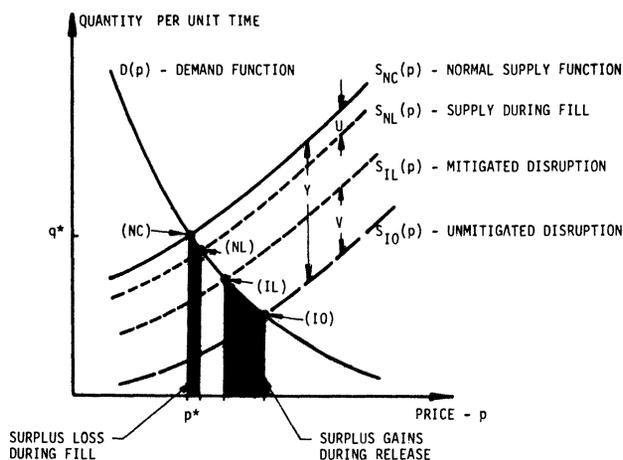


FIGURE 2. Consumer Surplus Implication of a Stockpiling Policy.

reference and calculate the consumer surplus loss associated with each price level relative to p^* .

The expected consumer surplus loss due to supply insecurity is then given by the sum of the consumer surplus losses associated with each price level weighted by the probability of being in that supply state. In addition to the expected losses of consumer surplus, the cost of supply insecurity includes a holding cost which is assumed to be proportional to the average stockpile level \bar{L} . Part of these losses are offset by the appreciation in the stored oil which is purchased at price p_{NL} and released at price p_{IL} . The expected steady state net insecurity cost per unit time is thus given by

$$\begin{aligned}
 J(C, U, V) &= P_{IO}W(p_{IO}) + P_{IL}W(p_{IL}) + P_{NL}W(p_{NL}) && \text{Consumer Surplus Losses} \\
 &+ H\bar{L} && \text{Holding Cost} \\
 &- (P_{IL}Vp_{IL} - P_{NL}Up_{NL}) && \text{Stock Appreciation}
 \end{aligned}$$

where

$$W(p) = \int_{p^*}^p D_{US}(\xi) d\xi$$

and H is the holding cost per unit stock per unit time.¹

Note that without stockpiling, the expected steady state insecurity losses are given by

$$\hat{J} = (P_{IO} + P_{IL})W(p_{IO}) = (1 - s)W(p_{IO}).$$

The expression derived above for the expected insecurity cost contains essentially the same components as the single stage cost in the dynamic programming formulations mentioned earlier. Here, however, since we assume steady state, we seek an optimal steady state stockpiling policy which will minimize the expected insecurity loss rate $J(C, U, V)$. The attractive aspect of this approach is that it reduces the problem to a static optimization problem.

To obtain an explicit formula for the objective function $J(C, U, V)$ we assume a specific demand function of the form: $D(p) = \bar{q} + ap^{-\epsilon}$. This is the same function used by Teisberg (1981) except that he had an additional time dependent term e^{gt} multiplying the entire expression on the right. For simplicity we also assume that supply is inelastic and fixed at level q^* during a normal period and at $q^* - Y$ during an interruption. We then obtain the prices in the four states by equating supply and demand in each state and solving for p . For a normal supply rate q^* we have $q^* = \bar{q} + ap_{NC}^{-\epsilon}$ which yields

$$p_{NC} = a^{1/\epsilon}(q^* - \bar{q})^{-1/\epsilon} \triangleq p^*.$$

Repeating this calculation for each of the other three states and substituting the above expression for p^* yields the corresponding prices:

$$p_{IO} = p^*[1 - Y/(q^* - \bar{q})]^{-1/\epsilon}, \quad p_{NL} = p^*[1 - U/(q^* - \bar{q})]^{-1/\epsilon},$$

$$p_{IL} = p^*[1 - (Y - V)/(q^* - \bar{q})]^{-1/\epsilon}.$$

Next we calculate the consumer surplus losses in the various states. We obtain after

¹ It should be noted that while the prices are based on the total world oil demand function $D(p)$, the consumer surplus calculations are for the U.S. only and are therefore based on $D_{US}(p)$.

some manipulation

$$W(P_{IO}) = \gamma p^* \left[\bar{q} \left\{ \left[1 - Y/(q^* - \bar{q}) \right]^{-1/\epsilon} - 1 \right\} + (q^* - \bar{q}) \left\{ \left[1 - Y/(q^* - \bar{q}) \right]^{(\epsilon-1)/\epsilon} - 1 \right\} / (1 - \epsilon) \right].$$

The expressions for $W(p_{IL})$ and $W(p_{NL})$ can be easily obtained by replacing Y in the expression for $W(p_{IO})$ with $(Y - V)$ and with U respectively.

It is convenient at this point to normalize all the quantities through division by $(q^* - \bar{q})$. The normalized quantities will be denoted by lower case letters, i.e., $(u, v, y, l, c) = (1/(q^* - \bar{q}))(U, V, Y, \bar{L}, C)$. We also define $h = H/p^*$ and $F = J/p^*(q^* - \bar{q})$. Now, using the probabilities given in §2 and the expressions for the consumer surplus given above, we obtain the following formula for the normalized insecurity cost F

$$\begin{aligned} F = & \gamma(1-s) \left\{ \alpha \left[(1-y)^{-1/\epsilon} - 1 \right] + \left[(1-y)^{(\epsilon-1)/\epsilon} - 1 \right] / (1-\epsilon) \right\} \\ & + \gamma P_{NL} \left\{ \alpha \left[(1-u)^{-1/\epsilon} - 1 \right] + \left[(1-u)^{(\epsilon-1)/\epsilon} - 1 \right] / (1-\epsilon) \right. \\ & \quad \left. - (u/v) \alpha \left[(1-y)^{-1/\epsilon} - (1-y+v)^{-1/\epsilon} \right] \right. \\ & \quad \left. + (u/v) \left[(1-y)^{(\epsilon-1)/\epsilon} - (1-y+v)^{(\epsilon-1)/\epsilon} \right] / (1-\epsilon) \right\} \\ & + \gamma P_{NL} u \left[(1-u)^{-1/\epsilon} - (1-y+v)^{-1/\epsilon} \right] + hl, \quad \text{where} \\ & P_{NL} = s \{ 1 - (\rho - 1)e^w / (\rho - e^w) \}, \\ & l = c \{ 1 - [\rho - (1-s + \rho s)(e^w - 1)/w] / (\rho - e^w) \}, \\ & w = c(\mu/v - \lambda/u), \quad \rho = \lambda v / \mu u, \quad \alpha = \bar{q} / (q^* - \bar{q}). \end{aligned}$$

When λ , μ , h and y are given, F is determined by the stockpiling policy variables c , u and v . These variables express the reserve capacity, fillup and release rate in multiples of $(q^* - \bar{q})$.

The stationarity of the demand function and supply levels is perhaps the most controversial simplification in our model, since it leaves out the effects of economic growth and the exhaustibility of oil. These dynamic effects have been included in the various dynamic programming formulations. Teisberg (1981), for instance, accounts for economic growth by including a multiplicative exponential growth term in the demand function. He also adjusts the normal supply level over time so that the normal price follows an exponential growth trajectory at an annual rate of 2%. This price increase reflects the exhaustible nature of oil and is consistent with classic models (see, e.g., Hotelling 1931) of optimal exhaustible resource depletion. This price increase is also employed by Teisberg to justify a finite time horizon of 25 years, since at that point, he argues, price of oil will reach a level at which nonexhaustible sources of oil substitutes (backstop technologies) become available. Recent work by Powell (1983) and by Oren and Powell (1985) suggests that in the transition to backstop energy resources, the oil prices will "overshoot" above the predicted long-run backstop prices. Such an overshoot will be caused by learning curve effects and by the gradual (rather than instant) buildup of backstop production capacity. Unfortunately, the results of the dynamic programming formulation are quite sensitive to the assumed oil price at the transition time to backstop technologies, since this is the sell-off price of the reserve at the end point. This might suggest that a refined model accounting for the known

stationary aspects of demand and price is not worthwhile, unless a more elaborate model of the transition to backstop technologies is employed.

One clear advantage of our model is that its results do not depend on any sell-off price assumption. To understand, however, the other implications of the demand and supply stationarity we will describe now an alternative formulation resulting in an objective function that closely resembles the one used in our analysis. As in Teisberg's (1981) model we assume a demand function of the form $D(p, t) = e^{gt}(\bar{q} + ap^{-\epsilon})$, and a normal price trajectory of the form $p^* = p_0^* e^{dt}$. In addition, we assume that the holding cost H is proportional to p^* , i.e., grows exponentially at rate d . Furthermore, the shortfall quantity Y and the policy variables C, U are proportional to $(q^* - e^{gt}\bar{q})$ i.e., they grow exponentially at rate $(g - \epsilon d)$. Under these assumptions, the normalized quantities u, y, c, u, v and l are again constant, provided that we replace $(q^* - \bar{q})$ with $(q^* - e^{gt}\bar{q})$. The expected insecurity cost rate is then given by

$$J = p^*(q^* - e^{gt}\bar{q})F = ap_0^{*1-\epsilon} e^{[g+(1-\epsilon)d]t} F$$

where F is defined as before. The only time dependent term in the expression for F is α which is now given by

$$\alpha = e^{gt}\bar{q}/(q^* - e^{gt}\bar{q}) = (\bar{q}/ap_0^{*\epsilon})e^{\epsilon dt},$$

i.e., it grows exponentially at rate ϵd . Since, however, ϵ is of order 0.1 (see next section) we will neglect the time dependency of α and treat F as stationary. Minimizing F is then equivalent to minimizing a discounted present value of a cost stream J over any finite time horizon, with a discount rate $[g + (1 - \epsilon)d]$. We may conclude, therefore, that the normalized policy parameters determined by optimizing our stationary model are also "near optimal" with respect to the discounted present value of the expected insecurity cost in the less restrictive model when the discount rate offsets the combined growth rates of demand and price. That is almost the case in Teisberg's (1981) base case analysis, where g and d are assumed to be 2% each while the discount factor is taken as 5%. When the discount rate is higher than $[g + d(1 - \epsilon)]$, the above discussion suggests that the stationary analysis will produce optimal policies which put too much weight on future costs and, hence, will be more conservative than a fully dynamic optimization. The stationary analysis will, therefore, tend to overestimate the optimal reserve size and fill rate while underestimating the release rate.

4. Computational Results and Sensitivities

Using the objective function developed in the previous section we calculated the optimal policy variables u, v and c and the reduction in the insecurity losses, for a variety of assumptions. In particular, we considered various values of the parameters λ, μ, y, h , and ϵ which characterize the frequency and duration of supply interruptions, their magnitude, the holding cost and the demand elasticity. The numerical optimization was performed using a quasi-Newton algorithm described by Gill, Murray and Pitfield (1972). This algorithm is particularly suitable for our complicated objective function since it requires only function values and uses finite differences approximations of the derivatives. Another attractive feature of the program is the local random search procedure at the convergence point, the purpose of which is to avoid convergence to a local minimum.

We describe now the input data for our computational analysis and the motivation for various assumptions. Some of the numbers we have used may seem out of date but we used them so that we can compare our results to those obtained in earlier studies.

TABLE 2
Supply Shortfall Assumptions

	Shortfall	Fraction Shortfall	Crude Oil Price/Barrel
Normal (30 mbd)	0	0	\$34.3
Minor Disruption	2 mbd	0.07	46.7
Moderate Disruption	6 mbd	0.20	89.1
Major Disruption	9 mbd	0.30	148.8

Demand Model Parameters

The base case parameters for the demand function were chosen to be $\bar{q} = -36$, $a = 94$, $\epsilon = 0.1$. Table 2 gives the oil prices implied by these parameter values at different levels of supply shortfall. The normal supply level was assumed to be $q^* = 11$ billion barrels/year or equivalently 30 million barrels per day (mbd), which was approximately the world traded oil in 1978. This assumption is common to all the SPR analyses referenced above. The normal price of imported oil is assumed to be \$34.3 which was the 1981 price (according to Weekly Petroleum Status Report, February 6, 1981). Chao and Manne also assume a normal price of \$34/barrel, while Teisberg (1981), Balas (1981) and Hogan (1982) assume a normal price of \$30/barrel in 1979 dollars (that is roughly equivalent to \$34/barrel in 1981 dollars).

The other prices in Table 2 are roughly consistent with projections given in a recent study of the World Oil Market published by the Energy Modeling Forum (1982), Stanford University. This study projects, for instance, a short-term increase of around \$100.00 per barrel for a 9 mbd shortfall. Such a major shortfall would correspond to a total interruption of Saudi Arabian oil production.

The demand elasticity of our assumed demand function at any demand level q is given by $\sigma = \epsilon(1 - (\bar{q}/q))$. Thus for $\epsilon = 0.1$, $\bar{q} = -36$ and $q = q^* = 30$ the local elasticity is $\sigma = 0.22$. On the other hand, over the range of 9 mbd shortfall the average elasticity is

$$\sigma = \frac{9}{30} / \frac{(148.8 - 34.3)}{34.3} = 0.09.$$

To explore the sensitivity of the results to demand elasticity we will keep the parameters a and \bar{q} of the demand function constant and vary ϵ . In particular we consider $\epsilon = 0.06, 0.08, 0.1, 0.15$. The fraction of traded world oil imported by the U.S. is assumed, as in Teisberg's (1981) paper, to be $\gamma = 0.25$ for $q^* = 30$ mbd. This means a normal U.S. oil import of 7.5 mbd.

Supply Model Parameters

We assume an average duration of supply interruption to be about 8 months and consider three alternative mean times between interruption: 5, 7.5 and 10 years. This implies $1/\mu = 0.67$ and $1/\lambda = 5, 7.5, 10$. We consider several levels of interruption described by the parameter ratios $Y/q^* = 0.1, 0.15, 0.2, 0.25, 0.3$. The maximum interruption corresponding to $Y = 0.3q^*$ implies a shortfall of 9 mbd, which is about the total Saudi oil production.

Unit Holding Cost

As in Teisberg (1981), we estimate the physical storage cost of holding one barrel of oil at \$1/year and assume a stockpile acquisition cost of \$38/barrel (which is about 10% higher than the normal international price of $p^* = \$34$ /barrel). Taking the annual cost of capital of stored oil to be 8%, we get an annual holding cost of $H = 1 + 0.08 \cdot 38 = \4 per barrel, and $h = H/p^* = 4/34 = 0.12$.

TABLE 3
State Probabilities, Prices and Consumption for Base Case

State	NC	NL	IL	IO
Probability	0.65	0.285	0.050	0.015
p/p^*	1.09	1	2.54	4.49
q/q^*	0.98	1	0.81	0.7

Base Case Results

As a base case we assume the parameter values $\epsilon = 0.1$, $1/\mu = 0.67$, $1/\lambda = 10$, $h = 0.12$, $Y/q^* = 0.3$, $q^* = 11$ billion barrels/year, and $\gamma = 0.25$. The resulting optimal policy variables for this case are as follows.

The stockpile acquisition rate U equals $0.019q^*$, which is equivalent to 0.6 mbd. The optimal drawdown rate during an interruption V equals $0.109q^*$ or 3.3 mbd. The optimal reserve capacity C is $0.143q^*$, which equals 1.57 billion barrels of oil. Consequently, the expected average stockpile level \bar{L} is $0.122q^*$, which equals 1.34 billion barrels of oil or about 6.5 months of normal U.S. oil imports.

We also calculated the steady state probabilities, price increases and consumption reduction for each of the four supply states as shown in Table 3. Notice (Table 3) that the optimal stockpiling policy reduces the probability of a 30 percent interruption from 6.5% to 1.5%. The optimal release policy increases consumption during an interruption from 70% to 81% of the normal level and reduces the price by nearly 50% (as compared to the unmitigated interruption price). The optimal fill rate is quite low, diverting about 2% of normal world traded oil or 8% of U.S. normal oil import to the stockpile. This causes a 9% price increase.

The expected annual insecurity losses to the U.S., J , with the optimal steady state stockpiling policy is $0.082 p^* q^*$. Since the U.S. normal oil import q_{US}^* is 25% of the total amount of traded oil on the world market q^* , we have

$$J = 0.082 p^* q^* = 0.328 p^* q_{US}^* = 11.55 q_{US}^*.$$

This implies an approximate \$11.15/barrel insecurity cost or 33% of the U.S. oil import expenditure. Note that the actual insecurity cost per barrel should be somewhat higher since the average import consumption is reduced by occasional interruptions. The error, however, is quite small since, even without stockpiling a 30% shortfall for eight months, every ten years implies no more than a 2% average shortfall.

For comparison we also calculated the insecurity losses without stockpiling, which comes out to be

$$\hat{J} = 0.107 p^* q^* = 0.428 p^* q_{US}^* = 14.55 q_{US}^*.$$

Thus, an optimal stockpiling policy will reduce the insecurity losses by 24%, which is about 10% of the total U.S. oil import expenditure.

Earlier in the paper we raised the issue of private stockpiles and asserted that if the government attempts to implement an SPR policy which is optimal for the economy, then there will be no incentive for private stockpiles by risk averse agents. To verify this assertion we need to compute the expected rate of speculative gains due to stockpiling under the optimal policy. These gains, which we denote by G , consist of the expected capital appreciation on sold stocks less the holding cost, i.e., $G = P_{IL} \cdot V(p_{IL} - p_{NL}) - H\bar{L}$. Substituting the base case results into the expression for G gives

$$G = [0.05 \cdot 0.109(2.54 - 1.09) - 0.12 \cdot 0.122] p^* q^* = -0.026 p^* q_{US}^* = -0.88 q_{US}^*.$$

This calculation shows that the optimal SPR policy will run at an expected speculative loss, or, in other words, will require a public subsidy amounting to 2.7% of the U.S.

import spending or \$0.88 per imported barrel of oil. Hence, private risk neutral agents which do not enjoy such subsidy will have no incentive to duplicate such a policy. An alternative strategy that might be adopted by private stockpiles is to purchase oil when prices are at normal level and hold it for the event that the public reserve is depleted. According to Table 3, the stationary probability for such an event occurring in a unit time interval under the optimal SPR policy is 0.015 and the corresponding price is $4.49p^*$. Subtracting the holding cost from the expected speculative profit rate yields again a net loss of $0.015 \cdot 3.49p^* - 0.12p^* = -0.068p^*$ per barrel of oil. This result could have been anticipated since if such an action was profitable, consumer welfare consideration would have made it even more attractive and it would have been part of the optimal SPR policy. The above discussion illustrates the general principle that a socially optimal SPR policy will eliminate the incentive for private stockpiling by a risk neutral agent.

To complete this discussion we note that under our base case assumption, and without an SPR policy, the stationary probability of a full shortage is 0.065. Hence, the expected speculative profit rate for the marginal stockpiled barrel of oil is at least $0.065 \cdot 3.49p^* - 0.12p^* = 0.11p^*$, i.e., 11% of normal price. In fact this is a lower bound on the speculative profit since the expected holding cost per unit sold is less than the unit holding cost. Consequently, in the absence of an SPR policy there would be incentives for private stockpiles which could capture part of the insecurity cost savings targeted by the SPR. Under perfect competition and risk neutrality one expects that the equilibrium private stockpiles will be such that the expected speculative profits are zero. The calculation of such an equilibrium, however, is out of the scope of this paper.

Parametric Analysis

We now explore the effect of the various parameters on the optimal stockpiling policy and on the U.S. supply insecurity cost. In particular, we illustrate the variations in the optimal reserve capacity, average stockpile, fill rate and release rate as a percentage of the normal U.S. annual oil import q_{US}^* . We also evaluate the insecurity cost as a percentage of the U.S. import expenditures $p^*q_{US}^*$ under normal supply, and the probability of stockpile depletion. These quantities are displayed below as they vary with each of the three parameters y , $1/\lambda$, ϵ while the other parameters remain constant at their base case values.

Sensitivity to Disruption Size. Figure 3 illustrates the effect of disruption size on the optimal stock level and expected stockpiling gains. We observe that the relative advantage of stockpiling increases with the size of the anticipated disruption. Specifically, the ratio J/\hat{J} , giving the relative insecurity premium with and without stockpiling, declines almost linearly with disruption size. If $Y/q^* = 0.3$, $J/\hat{J} = 0.64$, which means that stockpiling can significantly reduce expected GNP losses from an oil import disruption. On the other hand, if oil supply interruptions are not expected to exceed a 10% shortfall, it is hardly worthwhile to stockpile, at least from an economic point of view. Naturally, the size of the stockpile also increases sharply with the expected disruption, as shown in Figure 3B. This increase, however, is considerably steeper than the comparable increase in fillup and release rates shown in Figure 3C. This indicates that as the expected disruption increases, the optimal stockpile management policy becomes more conservative, stretching the fill and depletion of the reserve over a longer period of time. This mitigates the adverse effect on price during the fill-up and provides more protection against stockpile depletion during a disruption. The latter effect is shown in Figure 3D which displays the probability of stockpile depletion under the optimal policy. The probabilities of the other three possible states are quite insensitive to variations in disruption size.

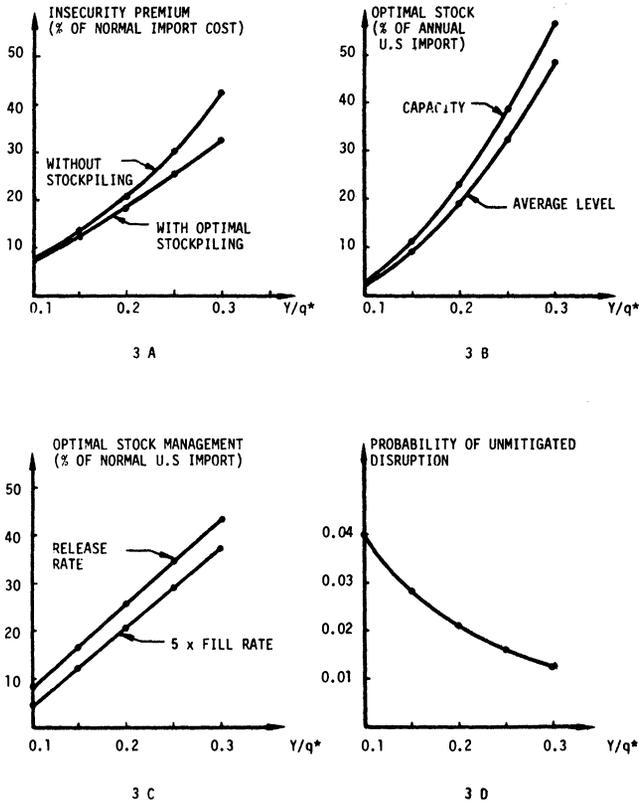


FIGURE 3. Sensitivities to Disruption Size— Y/q^* ($1/\mu = 0.67$, $1/\lambda = 10$, $\epsilon = 0.1$, $h = 0.12$).

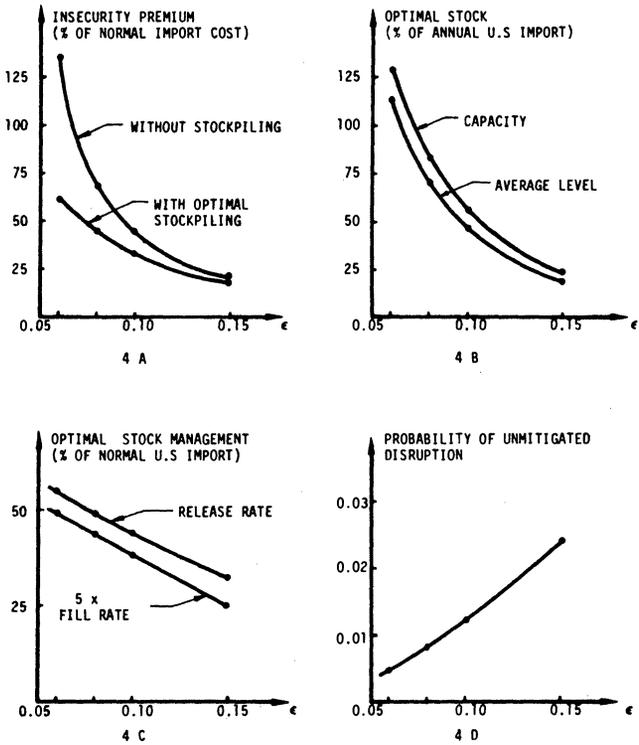


FIGURE 4. Sensitivities to the Elasticity Parameter— ϵ ($Y/q^* = 0.3$, $1/\mu = 0.67$, $1/\lambda = 10$, $h = 0.12$).

Finally we evaluated the expected subsidy requirements as described in the base case analysis for the different disruption sizes. This quantity appears to be nearly linear in disruption size ranging from 7¢ per imported barrel at $y = 0.1$ to 88¢ per imported barrel at $y = 0.3$.

Sensitivity to the Parameter ϵ . The parameter ϵ , which is directly related to demand elasticity, determines the consumer surplus losses due to price increases and has a strong effect on the optimal policy. As we see in Figure 4, a decrease in ϵ has a similar effect to an increase in disruption size. This could be anticipated since a lower elasticity will lead to higher prices due to supply shortfall, so the consumer surplus loss will be higher as if the disruption were bigger. We notice that for $\epsilon \sim 0.15$, stockpiling benefits are very small. If, however, ϵ is closer to 0.06, the gain from stockpiling is quite dramatic. Again we notice that when the optimal stockpile is larger, the fill up and release are spread over longer periods of time and the probability of total reserve depletion is reduced.

Sensitivity to Average Disruption Frequency. Figure 5 illustrates the effect of the parameter $1/\lambda$ representing the mean time between interruptions. As one would expect, the size of the stockpile and the benefit from stockpiling increase, the more frequent are the interruptions. One would expect, however, that in absence of an SPR, frequent interruptions would also motivate speculative private stockpiles, which would capture part of the SPR benefits.

In addition to the above sensitivities we examined the effect of variations in unit holding cost. The sensitivity of the results with respect to this parameter appears to be small. This suggests that the effect of discounting, which we neglected in our formulation, would also be insignificant.

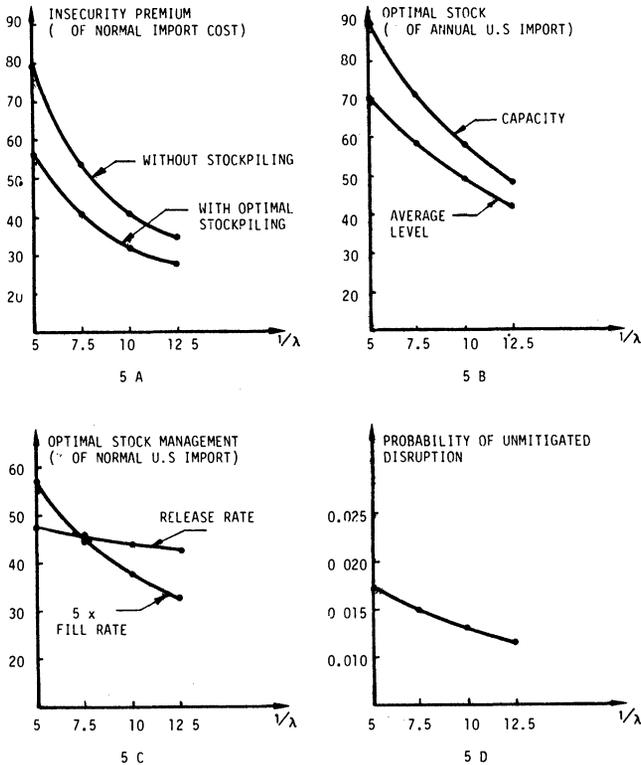


FIGURE 5. Sensitivities to Mean Time Between Disruptions— $1/\lambda$ ($Y/q^* = 0.3$, $1/\mu = 0.67$, $\epsilon = 0.1$, $h = 0.12$).

TABLE 4
Comparison with Other Models

Model	Optimal Stockpile Size	Acquisition Rate
Oren & Wan	1570 mb	0.6 mbd
Rowen & Weyant (1982) (Break-even analysis)	> 1500 mb	0.76 mbd
Balas (1981) (game theory)	833–2292 mb (8–10 month import)*	not specified
Teisberg (1981) (dynamic programming)	~1500 mb**	0.34 mbd (average) 0.82 mbd (initial)
Chao & Manne (dynamic programming)	1400 mb	1 mbd (empty) 0 mbd (full)

*Balas does not consider a single base case but rather a range of parameter values and recommends an SPR containing 8–10 months of import equivalent. The given range of the SPR size corresponds to a 7.5 mbd import rate.

**These results correspond to a “no interruption scenario” with potential shortfall of 20% for a year (our base case assumes 30% shortfall for 8 months). The given SPR size corresponds to a plateau level reached after 12 years at an average fill rate of 0.34 mbd. The initial fill rate during the first year is 0.892 mbd and it declines over time.

Finally, we examined the effect of using suboptimal reserve capacities while optimizing the fillup and release rates. It appears that the objective function is quite flat around the optimal C , indicating that deviations from the optimum will have only a small effect on the stockpiling benefits. This can be explained by the fact that the full reserve capacity will be seldom utilized.

Comparison with Results Based on Other Models

It is useful to compare our base case results with those obtained in the literature by other approaches. Unfortunately, different authors have used different base case assumptions in their analysis. To facilitate the comparison, therefore, we selected results corresponding to assumptions that are most compatible with our base case. The comparisons are summarized in Table 4.

It should be emphasized that since our results pertain to a steady state condition they do not take initial conditions into consideration. One expects that when the reserve is known to be empty the initial fillup rate should be faster than indicated by our analysis. The dynamic programming formulations are designed to address this issue and the results given in Table 4 suggest that the initial fill rate should be about 0.8 to 1 mbd.

5. Conclusions

The model presented in this paper attempts to fill the gap in the analysis of strategic oil reserve policies between simple two-period analyses and more elaborate dynamic programming approaches. The present approach captures some of the effects foregone by the two-period models at a significantly lower computational cost than the DP approach. The key simplification in our model is the assumption of fixed fillup and release rates. We also eliminate discounting although we account for interest on capital invested in the stockpile through our holding cost term. The latter approach is common in inventory theory in the derivation of The Economic Order Quantity, for instance. It is possible to generalize our model so that the state of the renewal process includes the continuous stockpile level while the fillup and release rates depend on that

level. The nonlinear programming problem then becomes an optimal control problem. We have foregone such a generalization in this paper since it significantly increases the computational cost and therefore defeats the purpose of our approach.

Comparison of our base case analysis with previously published results suggests that our approach is reasonably accurate and can provide a useful shortcut in evaluating the economic implications of strategic petroleum reserve policies and the optimal size of the SPR. Furthermore, the computational simplicity of this approach enables a wide range of sensitivity analyses, which are impossible with the dynamic programming formulations. Further analysis with our approach could include economic instruments such as tariffs and quotas as additional means for controlling the adverse impact of oil supply interruptions. This approach seems also promising for analyzing issues related to private stockpiles and their interaction with SPR policies. In particular we may use this model to evaluate the private stockpile level in the absence of an SPR and explore SPR policies which try to limit public subsidies by encouraging private stockpiling.

References

- BALAS, E., "The Strategic Petroleum Reserve: How Large Should It Be?," in *Energy Policy Planning*, B. A. Bayraktar et al. (Eds.), Plenum Press, New York, 1981.
- CHAO, H. P., "Inventory Policy and Supply Uncertainty," EPRI Working Paper, 1983.
- AND A. S. MANNE, "Oil Stockpiles and Import Reductions: A Dynamic Programming Approach," in *Energy Vulnerability*, J. Plummer (Ed.), Ballinger Publishing Company, Cambridge, Mass., 1982.
- CINLAR, E., *Introduction to Stochastic Processes*, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1975.
- Congressional Budget Office (CBO), "An Evaluation of the Strategic Petroleum Reserve," Report for the Subcommittee on Energy and Power, House of Representatives, Washington, D.C.; June 1980.
- Energy Modeling Forum, World Oil, EMF Study No. 6, Stanford University, 1982.
- Federal Energy Administration, Project Independence Report, Washington, D.C., November 1974.
- Gill, Murray & Pitfield, "The Implementation of Two Revised Quasi-Newton Algorithms for Unconstrained Optimization," National Physical Laboratory, Report No. DNAC11, 1972.
- HOGAN, W., "Import Management and Oil Emergencies," in *Energy and Security*, D. Deese and J. Nye (Eds.), Ballinger Publishing Company, Cambridge, Mass., 1981.
- , "Oil Stockpiling: Help Thy Neighbor," Harvard Energy Security Program, Discussion Paper H-82-02, March 1982.
- MEYER, R. R., M. H. ROTHKOPF AND S. A. SMITH, "Reliability and Inventory in a Production-Storage System," *Management Sci.*, 25, 8 (August 1979).
- OREN, S. S. AND S. G. POWELL, "Optimal Supply of a Depletable Resource with a Backstop Technology: Heal's Theorem Revisited," *Oper. Res.*, 33, 2 (1985), 277-292.
- POWELL, S. G., "The Transition to Nondepletable Energy," Ph.D. Dissertation, Engineering-Economic Systems, Stanford University, 1983.
- ROWEN, H. S. AND J. P. WEYANT, "Reducing the Economic Impacts of Oil Supply Interruptions: An International Perspective," *Energy J.*, 3, 1 (1982).
- TEISBERG, T. J., "A Dynamic Programming Model of the U.S. Strategic Petroleum Reserve," *Bell J. Economics*, 12, 2 (Autumn 1981).
- VARIAN, H. R., *Microeconomic Analysis*, Second Ed., W. W. Norton & Company, New York, 1984.