WP/12/278

IMF Working Paper

Modeling the Impact of Taxes on Petroleum Exploration and Development

James L. Smith

INTERNATIONAL MONETARY FUND

 WP/12/278

IMF Working Paper

Fiscal Affairs Department

Modeling the Impact of Taxes on Petroleum Exploration and Development

Prepared by James L. Smith*

Authorized for distribution by Michael Keen

November 2012

Abstract

We present a simple model of petroleum exploration and development that can be applied to study the performance of alternative tax systems and identify potential distortions. Although the model is a highly simplified, it incorporates many factors and some of the key tradeoffs that would influence an investor's investment behavior. The model recognizes the role of enhanced oil recovery and treats the impact of taxation on exploration and development in an integrated manner consistent with an investor's joint optimization of investments at both stages of the process. The model is simple and user-friendly, which facilitates application to a broad range of problems.

JEL Classification Numbers: H21, H25, L71, L72, Q32, Q38

Keywords: petroleum exploration, oil production, taxation, distortions

Author's E-Mail Address: jsmith@mail.cox.smu.edu

*This paper was prepared as a background paper for the IMF policy paper, *Fiscal Regimes for Extractive Industries: Design and Implementation* (http://www.imf.org/external/pp/longres.aspx?id=4701). The author is grateful to Philip Daniel, Ruud de Mooij, Michael Keen, and Alistair Watson for careful reading of a previous draft and many helpful comments. Any errors that remain are the responsibility of the author alone.

Contents Page

Tables

Figures

I. INTRODUCTION

This paper examines the impacts of taxation on petroleum exploration, development, and production. It focuses on generic versions of the fiscal regimes most commonly employed in petroleum-producing countries and examines the influence of those systems on the scope and efficiency of resource exploitation as well as the distribution of returns and risks between the investor and host government (HG). The analysis is facilitated by a new and relatively simple model of petroleum exploration and development that incorporates the way that taxes affect various margins of exploitation, including the extent of exploration in a given area, the timing and intensity of development of resulting discoveries, the timing and intensity of enhanced recovery methods, the overall recovery factor, and ultimate abandonment and closure of the field.

The investor's response to any given tax instrument or regime depends on the interrelations that link each of these decisions. For example, it may be supposed that a high royalty rate would cause early abandonment of a field and impair resource recovery. Holding all else equal, that is undoubtedly true. But a high royalty may also limit the intensity of the investor's initial development program, which might in turn cause production to decline at a slower pace, thereby extending the life of the field. In addition, however, the high royalty may discourage application of enhanced recovery methods as the field matures, and an investor who anticipates this may elect to increase investment in initial capacity as a more profitable alternative to enhanced oil recovery (EOR). The total impact of the royalty on resource recovery, the investor's rate of return, and government revenues depends on the solution to this set of interrelated investment problems.

The virtue of the model employed here is that it accounts for the investor's simultaneous consideration of all of these factors and describes the interrelated set of behavioral responses by which a rational investor would attempt to reduce the burden of whatever tax regime it may face. Those reactions, of course, are instrumental to the analysis of potential tax distortions and the regime's ability to capture resource rents for the HG. Because it is only an abstraction of what would otherwise be a highly complex optimization problem, the model is also quite simple and user-friendly, which facilitates analysis of a broad range of issues across a variety of circumstances. A general outline of the model is provided below.

A. Resource Development

The analysis is built on a new model of oil field development that integrates decisions regarding primary and enhanced recovery. The model is a direct extension of the traditional exponential decline model of oil field development often seen in the literature.[1] Initial investments to install productive capacity are taken with a view to the current oil price and

---

[1] See for example Smith and Paddock (1984), Adelman (1990), and Smith (1995a).

fiscal environment, but also with the knowledge that additional investments may be taken later on to enhance total recovery. The effect of investments taken to enhance recovery is to multiply the volume of remaining recoverable reserves by a fixed factor (determined by reservoir characteristics and technology). In the analysis reported below, we examine cases where the potential of EOR is to increase the recovery factor from roughly 33 percent in the primary phase of operations to roughly 45–55 percent overall.

Thus, at the development stage the private investor (hereafter “IOC” for International Oil Company) faces at least two decisions: (1) how much initial capacity to install, and (2) when, if ever, to commence enhanced recovery operations. We will also consider the IOC’s incentive to postpone initial field development, perhaps on expectations that prices will be higher in the future, or costs lower. Throughout the analysis we will focus primarily on differences in the way that alternative tax instruments and fiscal regimes impact these types of decisions.

The model of oil field development is simple, but designed to capture important operational tradeoffs that might be influenced by the method of taxation. For example, the margin between primary and enhanced recovery is resolved within the model. Investing in enhanced recovery too early is prohibitively expensive, but waiting too long will produce a relatively small increment to reserves. We assume the IOC has certain expectations regarding prices and technology and elects the development program that maximizes expected after-tax net present value (NPV).[2] Attributes of that program in terms of the size and timing of initial investment, the overall recovery factor, the initial extraction (decline) rate, the timing of enhanced recovery, and ultimate abandonment of the field are all determined within the model and recorded for each economic scenario (oil price level, cost level, field size) and fiscal regime under consideration (royalty, corporate income tax, various production sharing agreements, and resource rent tax).

The model is implemented and all optimizations performed within a relatively simple Excel spreadsheet. All cash flows (investments, costs, revenues, tax, and fiscal liabilities) are projected on an annual basis, although adapting the model to a quarterly schedule would be straightforward. For most of the results that follow, oil prices have been assumed to remain constant or follow a deterministic trend. In addition, we consider a mean-reverting random price process (calibrated to historical oil price movements) and examine how the respective fiscal regimes influence the distribution of risk between IOC and HG.

[2] The model is well behaved, meaning that the profit function is monotonic and a unique optimum always exists.

B. Resource Exploration

The oil field development model can be applied on a stand-alone basis to evaluate a previously discovered but undeveloped field. Or, it may be embedded in a larger model of exploration and discovery to study the impact of alternative fiscal regimes on exploration incentives and behavior. Exploration is assumed to consist of a series of exploratory wells, each with known cost. A well may produce a dry hole (non-commercial discovery), or one of three field types (small, medium, or large). For purposes of illustration, we assume those to hold 25, 100, and 750 million barrels of recoverable oil, respectively.[3] Discovery probabilities are derived from a physical model of geological information and drilling effectiveness. One factor is the explorationist's initial assessment that the block in question contains structures that have been geologically charged with hydrocarbons. We call this the "geologic" probability. The other factor is the conditional probability that any given exploratory well will succeed on the condition of the block having been charged with hydrocarbons. We call this the "technological" probability because, given the presence of hydrocarbons, the power and accuracy of the test performed by drilling an exploratory well is largely determined by technological factors. The implications of a geological probability of 70 percent coupled with a technological probability of 50 percent (and other combinations) will be illustrated.

An important implication of the law of conditional probability is that the probability of success on a second well declines after an initial failure, and so forth if the drilling sequence were to continue unsuccessfully.[4] If the explorationist's belief in the resource base was sketchy to begin with (low geological probability), this decline may be large. However, if the power of the drilling test is weak (low technological probability), the decline will not be so large since the failure may be attributed more to the failure of that specific well, not of the geological model underlying the area.

Within the model of exploration, the IOC is assumed to consider drilling a sequence of such exploratory wells, wherein discovery probabilities are updated after each failure. Absent success, the area will ultimately be abandoned (condemned) when the value of prospective discoveries (determined by the oil field development module) no longer justifies the increasing risk. Through its impact on the value of developed fields, the fiscal regime will therefore influence the extent of exploration—leading to earlier condemnation and less thorough search if the levies on production are too harsh. In addition, the fiscal regime may impact exploration directly through the operation of ring-fence provisions, which preclude

---

[3] These figures represent the maximum reserves that could be recovered via primary recovery operations. Additional volumes are available via enhanced recovery, but adverse fiscal provisions could prevent even the full volume of primary reserves from being produced.

[4] A technical derivation of this dynamic model of the probability of success is described in Smith (2005).

the IOC from using exploratory expenditures to defer tax liabilities from existing projects. Under a ring fence, expenditures on the first exploratory well, and any ensuing exploratory wells, must be carried forward to be weighed against future income and tax liabilities on the block in question if one of the exploratory wells is ultimately successful. Thus, the ring-fence provision has the effect of increasing the after-tax cost of exploration and poses a disincentive for thorough search. The strength of this effect will depend, however, on the other provisions of the fiscal regime that apply to developed fields. The potential magnitude of fiscal impacts on both exploration and development is documented within the study.

II. Related Research

Many papers have examined the economic impact of alternative systems of taxing extractive resources. An overview of that literature can be found in Smith (2012). Here we provide a brief summary that is intended to show primarily how modeling approaches differ, where gaps exist, and how the model developed in the present paper contributes to our ability to evaluate tax policy.

The performance of any system of resource taxation depends on (1) the ability to raise government revenue, (2) potential distortions of private investment that impair resource value, and (3) the resulting allocation of risk between government and investor. To fairly assess these factors, one must recognize the many ways by which informed taxpayers may alter their behavior to mitigate the tax.[5] A behavioral model is required, one that captures the potential for tax avoidance within the limits of the law and subject to the physical and economic constraints that define the extractive enterprise. Along these two dimensions (incorporating potential tax avoidance and accounting for the extractive nature of production) existing models vary significantly.[6]

Many tax studies are based on scenarios; the analysis begins with a description of an exploration and/or development program applied to a specific petroleum prospect ("model field"). This description comprises a fixed scenario that defines the volume of reserves, scale and timing of investment, number of wells, drilling success rates, intensity of development, initial production rate (and subsequent decline), variable operating costs, etc. One strength of the scenario approach is that these attributes can be patterned directly on industry experience for the region in question. The scenario, plus assumptions about the price of oil and relevant taxes, is sufficient to calculate expected cash flows and tax liabilities over the life of the

[5] As Triest (1998, p. 761) observed, "Reliable estimates of how tax incentives affect behavior are an essential input to the formation of tax policy." Similar sentiments can be found in Conrad and Hool (1984) and Poterba (2010).

[6] We limit discussion to the taxation of petroleum. Many of the approaches described here have also been applied to assess tax impacts on mining operations, as discussed in Smith (2012).

project and to determine the investor's return and government take. After repeating such calculations under alternative tax regimes, results are then scored to identify differences between regimes. The explicit projection of cash flows permits detailed modeling of even the most complex fiscal regimes. Examples include Kemp (1987, 1992, 1994); Van Meurs (1988, 2012); Smith (1995b, 1997); Schiozer and Suslick (2003); Johnston (2003); Tordo (2007); Johnston, Johnston, and Rogers (2008); and Daniel and others (2010).

Since much of the scenario is fixed a priori, however, the scope for potential tax avoidance behavior is limited. In effect, only the investor's initial decision to undertake the project (if after-tax cash flows meet the break-even conditions) or terminate production (when marginal after-tax returns become negative) are free to vary among regimes. Thus, it is possible to estimate how a given regime will affect the break-even price required for investment, or the minimum economic field size (assuming there are economies of scale), or the minimum required cost of capital, or the terminal flow rate that would trigger abandonment. It is not possible, however, to gauge the effect of the tax system on the intensity of initial development, the speed of production and/or subsequent decline rate, or the timing and magnitude of any secondary investments undertaken to enhance recovery—because these factors are all pre-determined. Any conclusions regarding the impact of a given tax regime on government revenue, resource rents, or the distribution of risk between parties are therefore subject to qualification.

Many researchers have relaxed the rigidity of the scenario approach—either by introducing flexibility with respect to the scale or timing of investment, but usually not both. Thus, we have few models that are designed to examine a taxpayer's comprehensive reaction to a given fiscal regime, or to assess the performance of real-world tax regimes that have the potential to influence multiple dimensions of a typical project.

Lund (1992), for example, allows the scale of development to adjust optimally to differences between tax regimes. The timing of investment and the shape of the production profile, however, remain fixed a priori. Compared to the rigid-scenario approach, allowing scale adjustments contributes significantly to the potential for distortion. Lund finds that incentives created by the Norwegian (1980) tax regime, for example, potentially reduce both the scale of development (by 50 percent) and total resource rents (by 25 percent). Blake and Roberts (2006) apply Lund's approach to five additional international regimes (Papua New Guinea, Alberta, Tanzania, Trinidad, and the São Tomé and Príncipe/Nigerian Joint Development Zone) and also find that recognizing the feedback effect of taxes on project scale increases the perceived potential for distortion.

Unlike Lund or Blake and Roberts, Zhang (1997) relaxes the development scenario via flexible timing. Here the investor is assumed to optimally delay the initial (and irreversible) investment in resource development as provided by the Real Options Theory. Again, the additional flexibility increases the scope of potential tax distortions, and Zhang shows (in the context of the UK Petroleum Revenue Tax regime) that a unique investment "uplift" is required to achieve neutrality with respect to timing. But, despite the flexibility of timing allowed in Zhang's analysis, the scale of development is assumed to remain fixed. As in the more rigid scenario approach, the scale of investment and size and shape of the resulting production profile are assumed to be determined independently of the tax regime. Panteghini (2005) also employs the Real Options Theory to study potential tax distortions of the timing and sequence of investments in an extractive enterprise. His model recognizes that an investor may elect to delay all, or just a portion, of the initial investment. But again, the scale of the total investment, the rate of extraction, and the amount of resource to be recovered are fixed exogenously.

Hyde and Markusen (1982) and Campbell and Lindner (1985) focus on the role of exploration in reducing geological uncertainty and devise models that identify potential tax distortions during the exploration phase of operations. Their models, like that in the present paper, treat the after-tax value of developed reserves as the primary incentive for exploration and recognize that the deductibility of exploration costs plays a key role. The scale of development that drives exploration, however, is again assumed to be determined independently of the tax regime in question.

Other researchers have attempted to replace the rigid production schedules characteristic of the scenario approach with simple reservoir simulation models that determine initial production and subsequent decline in accordance with the laws of physics. Jacoby and Smith (1985), for example, account for the impact of resource depletion on reservoir pressure and subsequent decline rates to determine the optimal rate of extraction and scope of investment. They find that royalties, windfall profits taxes, and net profit sharing regimes (as implemented in the United States during the 1970s) distort the minimum field size for commercial development, as well as the rate of extraction. They do not incorporate impacts at the exploration stage, however, and the simple "physics" embedded in the model are suited to the production of gas but not oil.

Uhler (1979) employs a more elaborate physical model for the production of oil, again based on the effect of depletion on reservoir pressure, and examines the impact of simple royalties and income taxes on the incentive for enhanced oil recovery via pressure maintenance. His model also integrates development and exploration activities by looking at the combined return from sequential investments in both—similar to what is done in the present paper. Helmi-Oskoui and others (1992) mount an even more elaborate attempt to join engineering-based reservoir simulation with economic analysis of optimal investment behavior. They assume, unlike Jacoby and Smith or Uhler, that the reservoir is physically heterogeneous (e.g., spatial variation in porosity and permeability), which adds many dimensions to the optimization problem since the location of wells, and not just their number, must be taken into account. The value of the added complexity seems doubtful, however, because simulations of the model produce implausibly low rates of extraction.

In summary, there have been many important contributions to our understanding of how taxes impact the various margins of resource development. The literature includes a wide variety of techniques and perspectives, all of which attempt to account for the taxpayer's behavioral response—if only to a minimal extent. As Poterba (2010) observes, understanding the behavioral response is what economic analysis adds to the accounting discussion of tax policy. The ideal approach would be to attach a richly detailed financial model of the fiscal regime to a robust production model of the extractive enterprise. As things stand, however, the more robust models of extraction are limited in their ability to accommodate complex fiscal structures of the sort encountered in the real world. The approach presented in the next section is an attempt to fill that gap. It provides an additional tool to examine how an extractive enterprise would adjust the intensity of exploration, the timing and intensity of initial development, the timing and intensity of enhanced recovery, and eventual abandonment of the field when facing real-world tax regimes.

III. The Modeling Approach

This section describes the structure of the model used to generate results. It also explains how exploration and development are integrated to provide a "full-cycle" analysis of investments and returns. To begin, we focus on a reservoir where the volume of original oil-in-place is fixed and denoted OIP.

A. Primary Production

During the primary phase of production, the rate of production is determined by naturally occurring conditions within the reservoir. Depending on the permeability of reservoir rock, the viscosity of trapped oil, and pressure gradients within the formation, oil is expelled by natural forces, and a certain fraction of original oil-in-place can be recovered without resorting to artificial stimulation. According to Total (2009), primary recovery amounts on average to about 33 percent of the original oil-in-place, although there is considerable variation among reservoirs. As a benchmark, therefore, we will assume that primary reserves are equal to one-third of original oil-in-place: $R = OIP/3$.

During the primary phase of production, output is assumed to decline from the initial level at a fixed rate over time:

$$Q_t = Q_0 e^{-at} \quad t \geq 0. \tag{1}$$

As discussed by Uhler (1979), this pattern results from the continuing loss of reservoir pressure as the oil is physically depleted. As will be seen, the rate of depletion, and therefore also the rate of decline, is determined by the intensity of the investor's development effort (e.g., the number of wells drilled into the formation).

By definition, the volume of primary reserves is the integral of equation (1):

$$R \equiv \int_0^\infty Q_t \, dt = \frac{Q_0}{a}. \tag{2}$$

It follows from equation (2) that $Q_0 = a \times R$, which means the decline rate (a) and the rate of extraction are the same. It also follows from (2) that the volume of reserves remaining in the reservoir at time *t* is given by:

$$R_t = \frac{Q_t}{a} = Re^{-at}. \tag{3}$$

The optimal intensity of development is determined by balancing the benefits of faster extraction against the cost of the required investment. Initial capital investment is assumed to depend on the size of field (*R*), the intended rate of extraction (*a*), and local conditions as reflected in a regional calibration factor (*A*):

$$I(a) = ARa^{1.68}, \tag{4}$$

where the elasticity of investment with respect to the rate of extraction (1.68) is taken from Smith and Paddock (1984). Since equation (4) is linear in *R*, constant returns to scale with respect to field size is implied, which is also consistent with Smith and Paddock's estimates, at least for onshore fields.

Investment requirements may be expressed alternatively in terms of the "capital coefficient," i.e., the amount of investment required per barrel of initial production:

$$CC(a, A) \equiv \frac{I(a)}{Q_0/365} = 365Aa^{0.68}. \tag{5}$$

This equation is calibrated to local conditions by solving equation (5) for *A*:

$$A = \frac{CC(a,A)}{365a^{0.68}}. \tag{6}$$

Experience with recent projects cited by *Petroleum Intelligence Weekly* (2009) indicates that capital coefficients range between $17,500 and $50,000 per initial daily barrel in the Middle East, but even greater variation can be expected elsewhere, such as in ultra-deep water or arctic conditions, for example. This variation enters the analysis through the parameter *A*.

B. Enhanced Production

At time *T*, the investor may elect to access additional reserves (enhanced recovery); remaining reserves are then augmented by the factor λ (>1):

$$Q_t^e = \lambda Q_t \; for \; t \geq T. \tag{7}$$

Notice that this specification assumes that the same decline rate will continue to apply to the expanded reserve volume. (This could be easily generalized, but would add one parameter to the optimization and require a more complicated optimization procedure.)

The state of reservoir depletion at time t is measured by the ratio of remaining reserves to the initial volume:

$$d_t = 1 - \frac{R_t}{R} = 1 - e^{-at}, \tag{8}$$

which increases from zero to 1 as extraction proceeds. Capital investment required for enhanced oil recovery (EOR) depends on the volume of remaining reserves to which EOR is applied (λR_T) and the state of depletion (d_T) at the time the EOR program is initiated:

$$I(\lambda, T, a) = \frac{Aa^{1.68}\lambda R_T}{d_T}. \tag{9}$$

This is the same functional form that applied to investment in primary recovery, but scaled according to the state of depletion. The intuition is as follows: implementing EOR early in the life of the field adds disproportionately to cost since, instead of adapting wells that have already watered out, new injection wells would have to be drilled, etc. As the field ages, it proceeds to a state where more idle capital is available to facilitate the enhanced recovery effort, thus reducing the incremental investment that is required.[7]

C. Optimal Field Development

To illustrate, we assume for the present a constant price and fixed operating cost per barrel over the life of the field. Thus, in the following expression, P can be interpreted as the constant net price per barrel of oil produced. In our applications, P is replaced by time-varying price and operating cost levels. We also impose no tax or fiscal costs here. Those additional terms are also added to the profit function in our actual applications, according to the specific fiscal provisions being investigated.

$$\begin{aligned} Max(a,T)\colon \pi^e(a,T) &= \int_0^T PQ_t e^{-rt} dt - I(a) + \int_T^\infty PQ_t^e e^{-rt} dt - \frac{e^{-(a+r)T}}{1-e^{-aT}} \lambda I(a) \\ &= PR\frac{a}{a+r}\left[1 + (\lambda - 1)e^{-(a+r)T}\right] - I(a)\left[1 + \frac{e^{-(a+r)T}}{1-e^{-aT}}\lambda\right] \end{aligned} \tag{10}$$

[7] Another reason to delay enhanced recovery is that overstimulation (adding too much pressure too soon) may cause water in the underlying aquifer to break through the formation, leaving unrecoverable oil trapped behind.

Notice that $\pi^e(a,T) \underset{T\to\infty}{\longrightarrow} \pi(a)$, where $\pi^e(a,T)$ represents the NPV of a development program that entails initial extraction at rate "a" and enhanced production from time T, and where $\pi(a)$ represents the NPV of the same development program but without enhanced production.

D. Solution Method

With no taxes or fiscal load, $\pi^e(a,T)$ has a closed form solution with first order conditions that can be used to identify a unique maximum, as in Smith and Paddock (1984). Alternatively, Excel's "Solver" can be used since T is treated as a continuous variable and the profit function is smooth and well behaved.

When tax levies and fiscal provisions are added to the model to create a treatment case, the specification of cash flows must be rendered in discrete time since fiscal regimes are usually defined in terms of discrete cash flows (periodic tax liabilities and shields). Therefore, for the applications discussed in the text, the model is discretized in the form of an Excel spreadsheet to express the contingencies and conditionality of the fiscal regime. In place of the continuous flow of production, first period production is given by $q_1 = aR$, and production in each subsequent period is given by: $q_{t+1} = (1\text{-}a)q_t$ (except in period T when production is given by $q_T = \lambda(1\text{-}a)q_{T\text{-}1}$), until the field is finally and permanently abandoned in the year that after-tax net cash flow becomes negative. To find the optimum development program for a given treatment, we perform a simple grid search over the control variables (a,T) via Excel's built-in "DataTable" function. (Because T is discrete, not continuous, Excel's built-in "Solver" function won't work). We find the objective function to be smooth and concave, which produces a unique solution and simplifies the grid search procedure.

E. Modeling Price Volatility and Financial Risk

Once the development program has been adopted and production capacity installed, the investor's costs are sunk and will be subject to the vagaries of future price movements. For any given development program (a,T), the magnitude and incidence of financial risk can be examined by replacing the assumed deterministic price projection with a simulated price process that allows prices to vary randomly through time. We have adopted a "mean reversion" model of random price movements for this purpose. (A pure random walk seems too extreme and unrealistic since many oil price movements are believed to be caused by transient supply and demand shocks rather than permanent shifts). Through Monte Carlo analysis, we then compute the mean and variance of NPV to the investor and to the HG for each fiscal regime. To be clear, this analysis is only used to measure the financial risk associated with the initial plan of development; it is not used to identify a plan of development that optimizes that risk.

It is assumed that periodic price movements follow a random process of the type:

$$P_{t+1} = P_t e^{s(lnK_{t+1} - lnP_t)} e^{u_{t+1}} \ \ where\ u_{t+1} \sim N(0, \sigma^2), \tag{11}$$

where s is the speed of adjustment to K_t, the long-term price trend. Equivalently, in terms of annual price changes:

$$\ln(P_{t+1}/P_t) = s(lnK_{t+1} - lnP_t) + u_{t+1} \tag{12}$$

Thus, the price tends to drop if it is above long-term trend and rise if it is below long-term trend. This random price process is easy to implement within the Excel-based discrete simulation model, and the speed of adjustment can be varied to produce more random-walk-like behavior if desired (zero adjustment speed = pure random walk).

F. Exploration

The investor is assumed to hold the right to drill a sequence of exploratory wells in a given block. The number actually drilled will depend, of course, on marginal costs and benefits. Each well is assumed to cost X, of which a fixed percentage (δ) represents intangible costs (relevant for tax accounting). In the absence of a ring-fence provision, intangible exploration costs can be expensed immediately to reduce income tax liabilities from other blocks (but not other types of tax levies—we assume strict ring fencing of production share and RRT regimes). If we let *RF* be a 0/1 indicator variable showing whether a ring fence is present, then after-tax cash flow for the period in which an exploratory well is drilled is given by the expression: $-X*[1-\delta*CIT*(1-RF)]$, where *CIT* represents the marginal corporate tax rate. Tangible exploration costs (or in the case of ring fence, all exploration costs) must be carried forward and can only offset income from the same block in the event a commercial discovery is made. Thus, total exploration costs carried forward after a series of n wells has been drilled is given by $CF(n) = n*X*[1-\delta\ (1-RF)]$—an amount that is added to other "recoverable" costs incurred during field development for purposes of computing the investor's future tax liabilities.

Success of each exploratory well is predicted by a physical discovery model in which there are four possible outcomes:

Small Field: $R = R_1$
Medium Field: $R = R_2$
Large Field: $R = R_3$
Dry Hole: $R = 0$

Drilling is assumed to continue until a discovery is made or the investor gives up, whichever comes first. The probability of outcome *i* from well *j* is denoted p_i^j, and is determined according to the discovery model in Smith (2005), which can be summarized as follows.

Let α represent the conditional probability that any given exploratory well will find a commercial field given that the investor's block is charged with hydrocarbons. And let β represent the probability that the block has been charged with hydrocarbons. It can then be shown that the dry hole risk of the n^{th} well in the sequence is given by: $p_4^n = 1 - \frac{\alpha(1-\alpha)^{n-1}\beta}{(1-\alpha)^{n-1}\beta+(1-\beta)}$ for *n* = *1, 2*, etc. If the relative likelihood of the three commercial field types are given by q_1, q_2, and q_3, the complete set of discovery probabilities is then established:[8] $p_i^n = (1-p_4^n)q_i$, for *i* = *1, 2, 3*.

The expected NPV of any individual well in the sequence (denoted V^n), taking account of dry hole risk and the range of possible discovery sizes, can now be calculated:

$$V^n = \sum_{i=1}^{3} \frac{p_i^n \times \Pi^e(a^*, T^*|R_i)}{(1+r)^{\Delta t}} - X[1 - \delta * CIT * (1 - RF)], \qquad (13)$$

… where the present value is computed relative to the date of drilling, the term $\Pi^e(a^*, T^*|R_i)$ refers to the optimized value of the given field size (from the field development module), *r* represents the investor's annual discount rate, and Δ*t* represents the time that lapses between exploration well and field development. Since the discovery probabilities (p_i^n) diminish with *n*, the NPV of an additional well eventually becomes negative and at that point (say it is after *N* wells have already been drilled) the investor would terminate exploration and abandon the block.

G. Integration of Exploration and Development

The full-cycle value of the complete exploratory sequence is given by the value of each of the *N* wells that constitute the exploration "campaign" ($V^n\ for\ n = 1, \dots, N$) multiplied by the probability that each of those wells gets drilled, denoted ϕ^n. Each of the *N* wells will get drilled if and only if all preceding wells in the sequence were dry. Thus: $\phi^n = \prod_{j=0}^{n-1} p_4^j$, where for convenience we have defined $p_4^0 = 1$. Thus, expected full-cycle NPV of the investor's right to exploit the block in question is given by:

$$NPV^{FC} = \sum_{j=1}^{N} \frac{V^j \phi^j}{(1+r)^{\Delta w}} \qquad (14)$$

… where the term Δ*w* represents the time (measured in years) that elapses between each well in the sequence.

[8] Formally, q_i represents the conditional probability of field type *i* given that a commercial discovery is made.

H. Fiscal Regimes Considered

Seven generic types of fiscal regime are considered:

1. NT: a "no-tax" benchmark
2. ROY: a simple royalty regime
3. CIT: a simple corporate income tax regime
4. PSC: a fixed-rate production-sharing contract
5. RF: an ascending R-factor production-sharing contract[9]
6. IRR: an ascending IRR-based production-sharing contract
7. RRT: a dual-tier resource rent tax.

Two versions of Regimes 4 and 5 are examined, with and without the option to apply interest on balances carried forward for purposes of cost recovery. Each regime is constructed of hypothetical parameters, schedules, and provisions. These regimes may be considered "pure" examples of the respective tax instruments. We do not consider hybrid regimes that consist of, say, royalties combined with an income tax, or income tax combined with production sharing, although these are common in practice.

The parameters of each fiscal regime are calibrated to capture two-thirds of the rent generated by an investor who operates subject to its provisions. Since the regimes have disparate effects on investment, this does not imply that they all raise the same amount of revenue for the government, but they are comparable in terms of the share of rents that accrue to the government. The analysis is not intended to replicate the specific circumstances in any particular country, but to highlight the characteristic effects of the individual instruments. Given the flexibility of the model, it would be straightforward to subject any particular real-world regime to a similar analysis.

More details of the chosen fiscal regimes are tabulated in Table 1. Additional background parameters used in the analysis are documented in Table 2.

[9] "R-factor" is the simple ratio of the investor's cumulative revenue to cumulative expense (undiscounted). Since expenditures precede production, the R-factor starts at zero and grows through time as the investment is recovered.

Table 1. Guide to Fiscal Regimes and Background Parameters Used in the Analysis

NT:	No-tax scenario. Only real economic costs are considered in the analysis. Government take is zero by definition. Useful as a benchmark of economically efficient resource exploitation.
ROY:	Fixed royalty (45 percent of gross sales value). No other government taxes or levies.
CIT:	Corporate income tax (61 percent with 15-year straight-line depreciation).
PSC:	Fixed production sharing contract, with cost oil limited to 60 percent of output, with government share of profit oil = 65 percent.
PSC+i:	Fixed production sharing contract as above, allowing interest on all balances carried forward for subsequent recovery. Government share of profit oil = 67 percent.
RF:	Production sharing contract with progressive R-factor. Government share starts at 46 percent but increases by increments of 5 percent (51, 56, 61, 66, and 71 percent) as the R-factor reaches 1.0, 1.1, 1.2, 1.3, and 1.4. Cost oil is limited to 60 percent of output.
RF+i:	Production sharing contract with progressive R-factor, allowing interest on all balances carried forward for subsequent recovery. Government share of profit oil starts at 52 percent but increases by increments of 5 percent (57, 62, 67, 72, and 77 percent as R-factor reaches 1.0, 1.1, 1.2, 1.3, and 1.4. Cost oil is limited to 60 percent of output.
IRR:	Production sharing contract based on investor's internal rate of return. Government share starts at 52.5 percent but increases by increments of 5 percent (57.5, 62.5, 67.5, 72.5, 77.5, 82.5, and 87.5 percent) as the IRR reaches 10, 15, 20, 25, 30, 35, and 40 percent. Cost oil is limited to 60 percent of output.
RRT:	Resource Rent Tax with dual tier threshold rates of return equal to 8 and 12 percent and tax rates of 50 and 67 percent. Thresholds are reset after each major capital investment (e.g., enhanced oil recovery).

Table 2. Background Parameters

Other background parameters in the cash flow analysis are set as follows, unless otherwise specified:

Real discount rate:	8 percent per annum
Initial oil price:	\$100 per barrel
Oil price appreciation:	zero percent real, per annum
Variable cost:	\$20 per barrel
Operating cost:	2 percent of cumulative capital expenditures, per year
Capex (development):	\$40,000 per initial daily barrel of production capacity (assuming 10 percent extraction rate)
Capex elasticity:	1.68 (elasticity of development cost with respect to extraction rate)
Field size:	100 million barrels recoverable via primary recovery (300 million barrels resource-in-place)
Enhanced Oil Recovery (λ):	2.5 (this enhancement factor measures augmentation of remaining recoverable reserves at the time EOR investment is made. (e.g., EOR = 2.5 means remaining recoverable reserves are multiplied by the factor 2.5)
Capex (exploration):	\$75 million, per exploratory well
Discovery probability:	50 percent = prob(successful well given presence of oil in block) = α
Geological probability:	70 percent = prob(oil present in block) = β
Interest on costs carried-forward:	zero percent per annum
Uplift for cost recovery:	zero percent
Ring fence:	precludes current expensing of intangible exploration expenses against other sources of taxable income
Tangible exploration expense:	20 percent of total exploration outlays

IV. OVERVIEW OF RESULTS

To illustrate the behavior of the model, we begin by describing the specific resource development programs that would be applied to a known discovery by the IOC under each fiscal regime, and given our benchmark economic assumptions. We also show how the presumed effectiveness of EOR ($\lambda = 2.0$ versus 2.5) affects that choice.[10] The general pattern is apparent in Figure 1, which plots the optimal extraction rate versus the timing of transition to EOR operations for the benchmark field (see previous Table 2 for values of the background parameters) under each fiscal regime. Under all fiscal regimes, the availability of more effective EOR technology ($\lambda = 2.5$) tends to reduce the initial rate of extraction (leaving more of the resource to be recovered via enhanced recovery) but accelerate the transition to enhanced recovery operations. The logic seems clear: both adjustments tend to expand the scope of the EOR project and subject a higher volume of reserves to the more effective enhanced recovery effort.

Figure 1. Impact of Enhanced Oil Recovery Effectiveness on Optimal Development

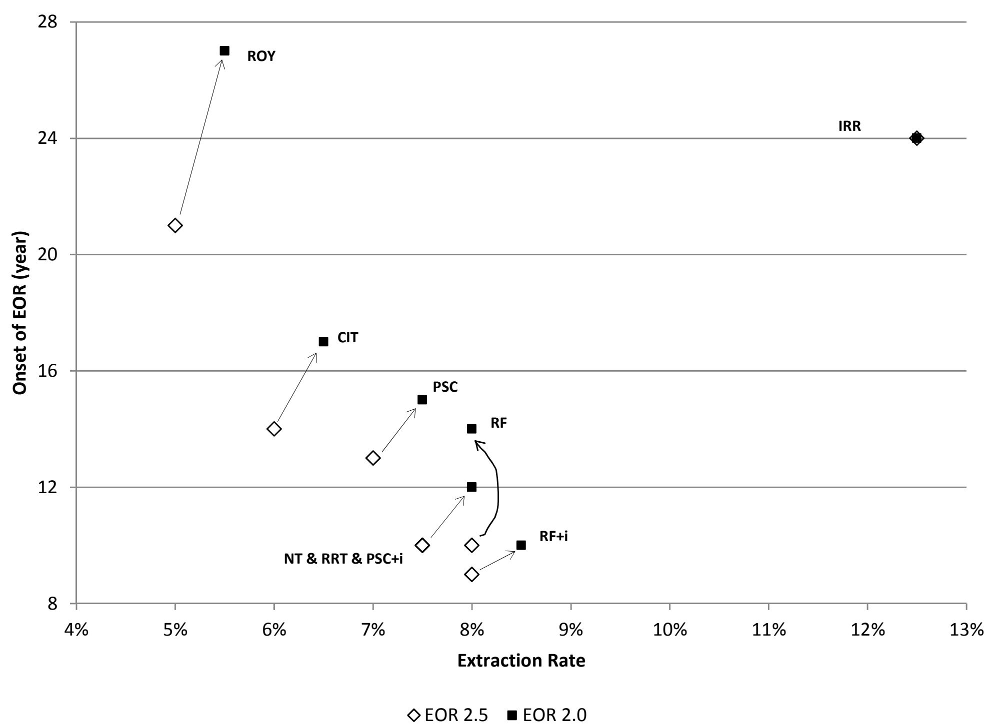

[10] Recall that the value $\lambda = 2.0$ means that EOR would double the volume of remaining reserves.

Figure 1 demonstrates the neutrality of the RRT and PSC+i regimes, which foster development as in the "no tax" case. The figure also reveals sharp differences between several of the other fiscal regimes. Under the NT regime (free of tax distortions), the IOC would extract initially at 7.5 percent per year and initiate EOR after 10 years of primary production. (Although the example is hypothetical, these numbers are not atypical of actual operations). The ROY regime (levied at 45 percent of gross sales value) markedly reduces the IOC's initial extraction rate because, left with a small fraction of the value of additional barrels, there is less incentive to incur the high cost of rapid extraction. Thus, the reservoir is depleted slowly (5 percent) and EOR is put off for many years. Under the PSC regime based on the IOC's cumulative IRR, EOR would not even be attempted, but the IOC would instead undertake very intense development (and high extraction rates) in the primary phase. Under the IRR-based regime, EOR is not feasible because by the time it would be initiated, the IOC has already achieved a high return, and that exposes the EOR program, which by nature is a marginal investment, to disproportionately high marginal tax rates. Zhang (1997) found a similar distortion within the UK Petroleum Revenue Tax regime where uplift did not extend to secondary investments. This shows the importance, in practice, of providing a mechanism to reset the production-sharing rate whenever major capital investments are undertaken.

The income tax regime also creates significant distortions that slow the pace of investment and extraction. The production-sharing regimes, in contrast, approach more closely the NT benchmark. Detailed effects of the respective fiscal regimes under high and low EOR effectiveness are summarized below in Table 3.

Table 3. Impact of Enhanced Oil Recovery on Resource Development and Recovery

	All Cases Shown Below Assume P= $100							
	λ = 2.0				λ = 2.5			
Fiscal	**Extraction**	**EOR onset**	**Recovery**	**Field Life**	**Extraction**	**EOR onset**	**Recovery**	**Field Life**
NT	8.0%	12	45%	43	7.5%	10	55%	44
Roy	5.5%	27	39%	60	5.0%	21	48%	64
CIT	6.5%	17	43%	58	6.0%	14	54%	61
PSC	7.5%	15	42%	44	7.0%	13	51%	46
PSC+i	8.0%	12	44%	38	7.5%	10	54%	39
RF	8.0%	14	42%	40	8.0%	10	53%	37
RF+i	8.5%	10	45%	33	8.0%	9	54%	35
IRR	12.5%	24	32%	23	12.5%	24	32%	23
RRT	8.0%	12	45%	43	7.5%	10	55%	44
Average	8.1%	16	42%	42	7.7%	13	51%	44

Figure 2 illustrates in economic and physical terms the contribution of EOR. Overall resource recovery factors increase from 33 percent if the IOC were restricted to primary recovery to well over 50 percent under EOR. This aspect of the model agrees with Total's (2009) reported experience with EOR, which typically delivers recovery of 50 percent or more. The

Figure 2. Impact of Enhanced Oil Recovery on Resource Recovery and Value

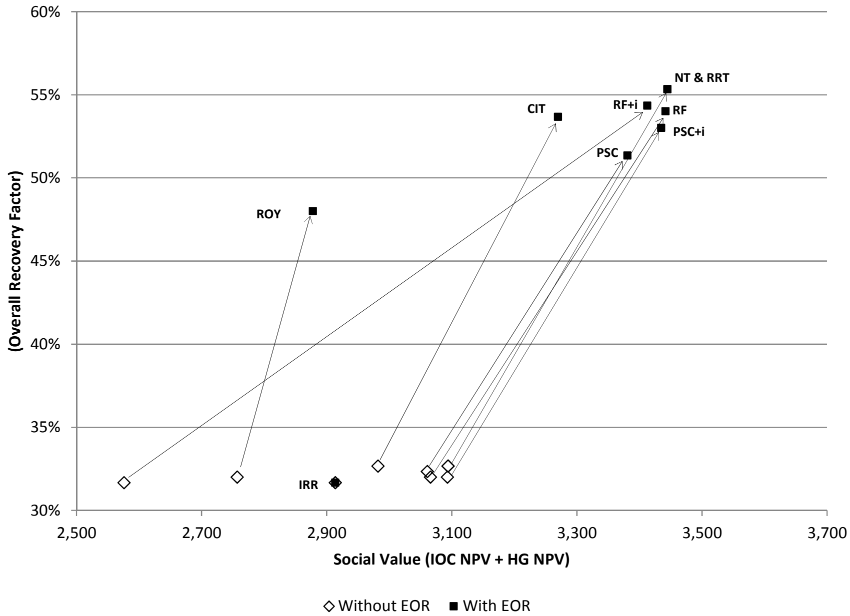

extra resource recovery adds significantly to the total profit available to be divided between the investor (IOC) and HG: on average adding 11 percent to the NPV of total cash flows. The only regime that deviates from the pattern is IRR, where high marginal tax rates preclude EOR. This underscores the importance of taking the IOC's decisions regarding EOR into account in any analysis of alternative fiscal regimes.

Figure 3 charts resource recovery by phase of production. While the IRR regime eliminates the EOR phase altogether, the ROY regime tends to maximize production during that phase. All of the other regimes capture more oil during the enhanced recovery phase than during primary recovery. Despite the higher volumes, however, the enhanced recovery adds less to overall NPV because it occurs later and is more expensive than primary recovery.

Figure 3. Resource Recovery, by Fiscal Regime and Phase

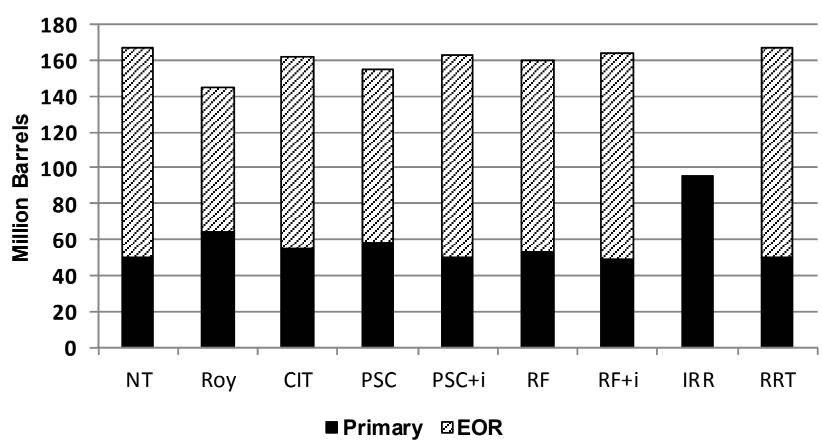

Finally, Figure 4 shows the overall impact of the respective fiscal regimes on the value of the field, and the division of value between IOC and HG. The distortions mentioned previously significantly reduce the realized value of the resource, at least in the case of ROY, CIT, and IRR regimes. The RRT regime is neutral and therefore does not impact the efficiency of recovery or impair value. The production-sharing regimes (apart from IRR) perform almost as well in that regard.

Figure 4. Total Net Present Value, by Fiscal Regime

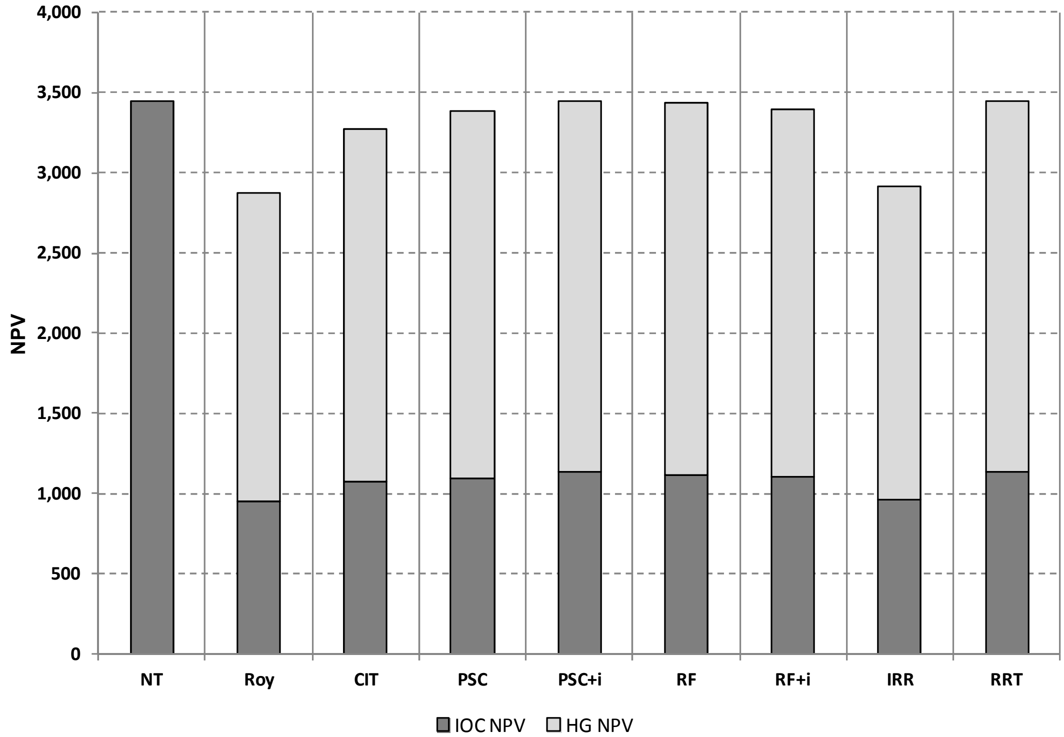

A. Intensity of Development

An increase in the price of oil may affect the IOC's choice of development program, but only if the fiscal regime does not capture all the extra revenue, which would leave the IOC cash flows unchanged. Considering the impact of a price increase from \$100 to \$150 barrel, Figure 5 shows the resulting impact on optimal extraction rate (primary recovery) and the transition to enhanced recovery. All of the tax regimes considered here direct some of the incremental revenue to the IOC (marginal tax rates are well below 100 percent, as we show later) and the result is more intensive development. Higher prices create strong incentives to get more oil out and get it out faster.

Figure 5. Optimal Development Programs, Price Impact

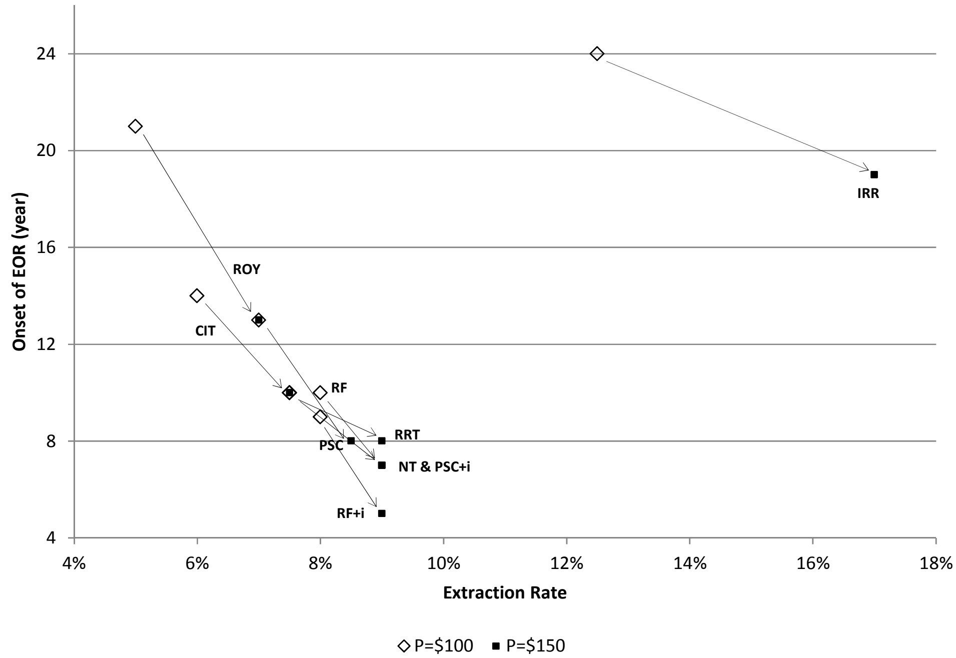

On average, the extraction rate increases from 7.7 percent to 9.4 percent, and the interval preceding enhanced recovery is reduced by more than 4 years. The efficient response to higher prices is indicated by the NT case, which shows extraction increasing from 7.5 percent to 9 percent and transition to EOR advancing by 3 years. It is clear from the figure that incentives created by the PSC and RF regimes trigger approximately the same response. However, the ROY, CIT, and IRR-based regimes are way off, regardless of the price level.

We assume the IOC will abandon the field once its after-tax net cash flow becomes negative, an event that occurs somewhere between 23 and 64 years after initial investment, depending upon the fiscal regime and price level (Table 3). At first glance, some of these results seem anomalous. For example, the fact that the relatively high royalty rate appears to extend the life of the field (cf. NT and ROY) conflicts with the usual presumption that royalties should hasten abandonment since royalties shrink revenues but not costs. That analysis is correct as far as it goes, but fails to incorporate the impact of the higher royalty on the IOC's development decisions. An important indirect effect of higher royalties is to reduce the scale of initial development, which by limiting the loss of natural reservoir pressure in the early years also tends to reduce the rate by which initial production declines.[11] The abandonment decision reported in Table 4 incorporates both factors, a higher royalty rate but also a higher late-life production rate.

---

[11] This is an inherent feature of the exponential decline model; reducing the rate of initial production also reduces the subsequent decline rate.

Table 4. Impact of Oil Price on Resource Development and Recovery

	All Cases Shown Below Assume λ = 2.5							
	P=$100				P=$150			
Fiscal	Extraction	EOR onset	Recovery	Field Life	Extraction	EOR onset	Recovery	Field Life
NT	7.5%	10	55%	44	9.0%	7	59%	38
ROY	5.0%	21	48%	64	7.0%	13	52%	48
CIT	6.0%	14	54%	61	7.5%	10	57%	51
PSC	7.0%	13	51%	46	8.5%	8	57%	38
PSC+i	7.5%	10	54%	39	9.0%	7	58%	34
RF	8.0%	10	53%	37	9.0%	7	58%	34
RF+i	8.0%	9	54%	35	9.0%	5	62%	29
IRR	12.5%	24	32%	23	17.0%	19	32%	16
RRT	7.5%	10	55%	44	9.0%	8	57%	40
Average	7.7%	13.4	50.7%	43.7	9.4%	9.3	54.8%	36.4

The impact of fiscal design on the transition from primary to enhanced production is shown in Figure 6, below. With price at $100 per barrel, the primary recovery phase ranges between 9 and 24 years, depending on fiscal regime, but lasts only 10 years in the NT case. Several of the fiscal regimes (ROY, CIT, and IRR) substantially retard the transition to enhanced recovery methods. Enhanced recovery itself typically continues on for another 30 years, although the high marginal tax rates imposed on enhanced recovery under the IRR-based PSC precludes EOR altogether, as we have noted previously. As is typical of large oil fields, commercial production extends over many decades, and would continue even longer than indicated if the real price of oil were assumed to rise during the life of the field—something that we introduce later.

Figure 6. Fiscal Impacts on Timing of Enhanced Oil Recovery and Abandonment

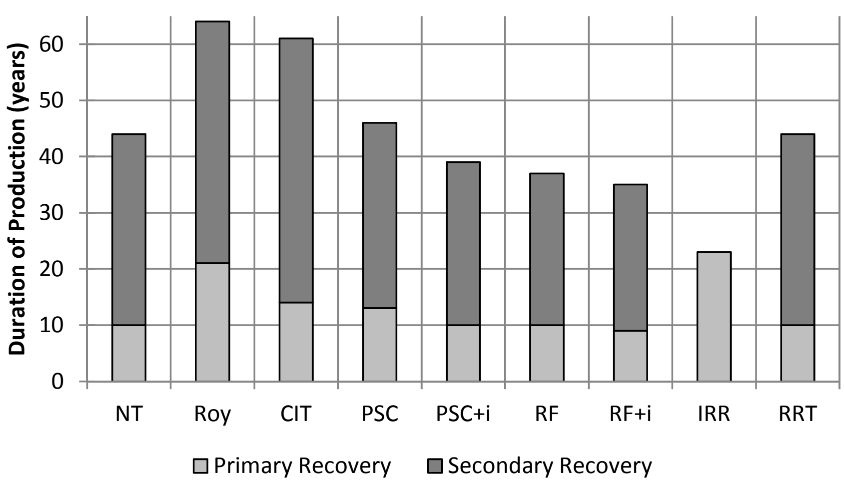

Abandonment under the RRT case occurs after 44 years, which corresponds to the NT case (no distortion). Abandonment occurs much earlier under the RF and IRR regimes due to the greater intensity of primary extraction they inspire. We focus on these timing effects only to illustrate the type of distortions that alternative fiscal designs impart on private investment decisions. Life of field is of no particular significance, per se; the shorter life obtained under some regimes might be better or worse (for the IOC, the HG, and for society as a whole) than the longer life obtained under other regimes, depending on the NPV of cash flows generated during the course of production. Under the circumstances assumed here, the longer life generated by the RRT regime happens to produce higher NPV for the IOC and society than any other regime, but it generates slightly less NPV for the government than the PSC+i and RF regimes.

B. Diligence

The previous calculations assumed that the IOC would commence development of the oil field immediately, but there are circumstances under which the IOC might wish to delay. Expectations of higher future oil prices, or lower future development costs, might cause a private investor to delay investments in order to take advantage of a more favorable environment. Of course, under these circumstances it could also be in the government's interest to postpone development since both parties might then be able to share in a larger pie. However, our particular interest here is to examine the possibility that certain ways of dividing the pie (i.e., fiscal regimes) might create false incentives or breed conflict between IOC and HG regarding the timing of development. The fact that many HGs apply pressure to IOCs to compel timely investment suggests that such conflicts do exist and perhaps are common.

The incentive to delay is highly situation-specific. Only if the NPV of the project were expected to rise by at least 8 percent per year (the assumed discount rate of the IOC) would the IOC benefit from delay. For projects with high profit margins (like those evaluated above), expectations of very rapidly rising prices, or very rapidly falling costs, would be required to justify a delay. But, for projects with a low per-barrel profit margin, even a modest rise in oil price could effect a substantial increase in IOC NPV, so delay becomes attractive. Thus, we should be concerned about diligence issues and conflicts mainly in the realm of high-cost, low-margin developments like oil sands, heavy oil, and deepwater or Arctic operations.

Figure 7 illustrates the strength of the private incentive to delay development of a high-cost field due to expectations of rising prices. For these calculations, most parameters are kept to previous values, but variable operating cost is increased to $60 per barrel (from $20); and although oil price starts out at the same $100 level, it is expected to appreciate in real terms at 3 percent per year. We have applied the development model described previously to calculate IOC NPV under each fiscal regime assuming the development (initial investment and all subsequent activities) is delayed by two and four years, respectively. In Figure 7, the resulting NPVs are reported relative to what the IOC would have earned without delay. Thus, it is evident from the chart that a four-year delay under the NT scenario would reduce the real economic value of the field by 5 percent, and development therefore should not be delayed if overall efficiency is the goal.

Figure 7. Incentive to Delay Development: High Cost Fields

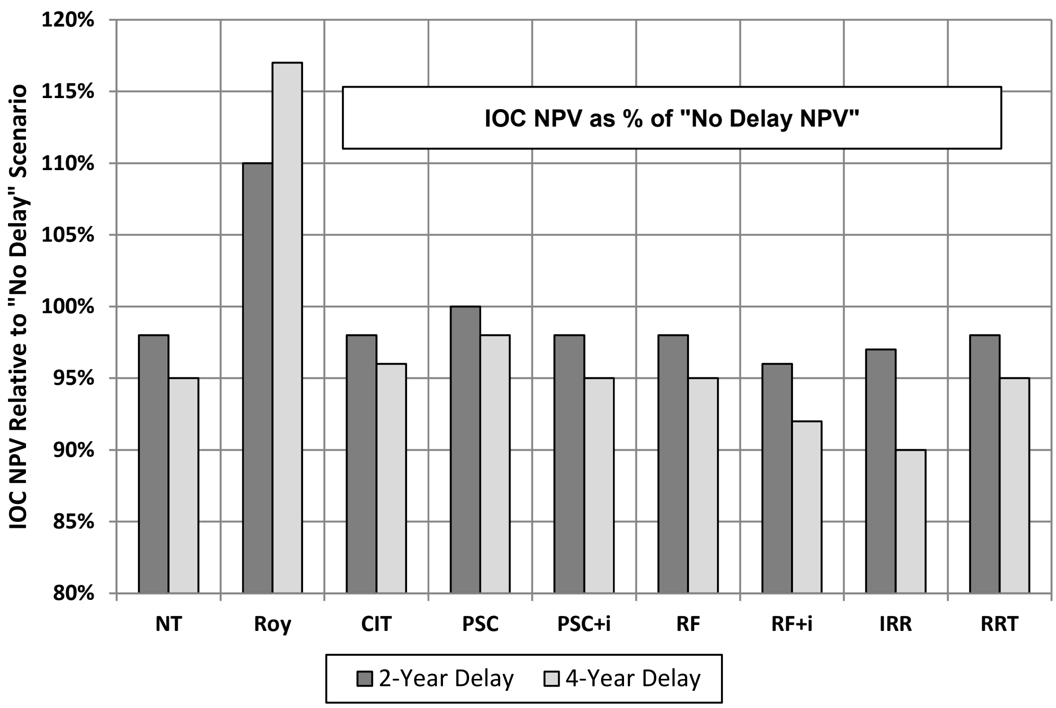

Notably, only the ROY regime creates an incentive for the IOC to delay development. Not only does the royalty reward delay, but the size of that reward grows with the length of the delay, which sets IOC and HG on a collision course. Their conflict is illustrated in Figure 8, which contrasts the impact of a four-year delay on IOC and HG, respectively. Whereas the IOC stands to gain, the HG (like society at large) stands to lose a significant portion (5 percent) of its total return if the schedule is left to the IOC. According to Figure 8, the interests of IOC and HG are aligned under all the other fiscal regimes, but it bears repeating that this is highly specific to the economics of the field in question. With still higher costs and thinner margins, conflicts might arise under other regimes as well. Clearly, this subject deserves more attention than we have given it, but if there is any general conclusion, it is that diligence issues will arise mainly in connection with regressive fiscal regimes, and only then in the case of high-cost, low-margin developments.

Figure 8. Royalties Create Timing Conflicts in High Cost Fields

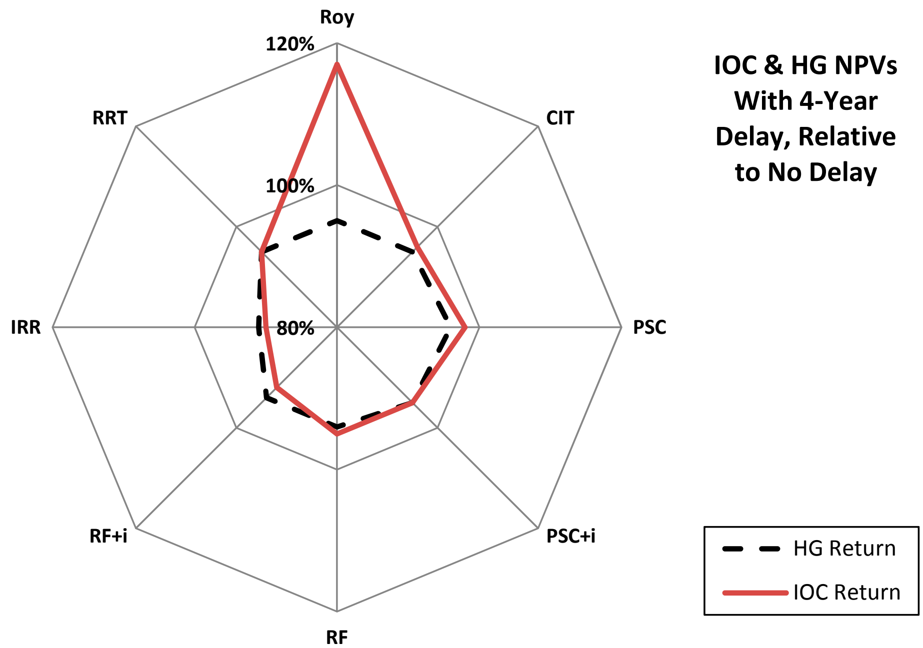

C. Fiscal Progressivity

A fiscal regime is said to be "progressive" if the government's share of total rents increases with the overall profitability of a project. Regimes that include tax instruments not closely linked to profits are more likely to be regressive, and the government's share will fall as oil prices rise. Progressivity may or may not be in the government's interest since it affects many factors, including the distribution of risk between IOC and HG, investment incentives, and public acceptance of financial results that are achieved over the entire course of the oil price cycle.

The first order of business is to identify which of the fiscal regimes chosen for study are progressive and which are not. Expectations would be that ROY regimes and flat-rate PSCs are regressive. To verify this hypothesis, however, it is necessary to describe in more detail the experiment that needs to be performed. Specifically, the question is whether or not, as prices change, all else is held constant. Do we envision the IOC holding to the same investment plan and production program under both low and high price scenarios? If those high prices are anticipated by the IOC, that result is unlikely. Therefore, should we instead consider that the IOC will react to the difference in price scenarios by revising its investment and production plans in ways that alter the level and timing of cash flows, and that may even confound the inherent progressivity of the fiscal regime? It seems the correct answer must depend on the context within which the question is asked, and we therefore perform both experiments.[12]

Figure 9 shows, for each fiscal regime, how government take is affected if the price of oil is \$150 instead of \$100 per barrel. Since all regimes are calibrated to capture 67 percent of total rents at the \$100 price, that 67 percent level is the relevant point of reference for judging progressivity. Two alternative measures of government take under the \$150 scenario are shown. The first (designated "Static IOC Response") assumes that, despite the higher \$150 price, the IOC keeps to its previous investment and production plan. This type of analysis would be appropriate for examining the consequences of price changes that are not anticipated. The second measure (designated "Dynamic IOC Response") assumes the IOC anticipates the \$150 price and adjusts its entire development program accordingly. The results obtained with these two measures are discussed below.

Figure 9. Government Take, by Fiscal Regime

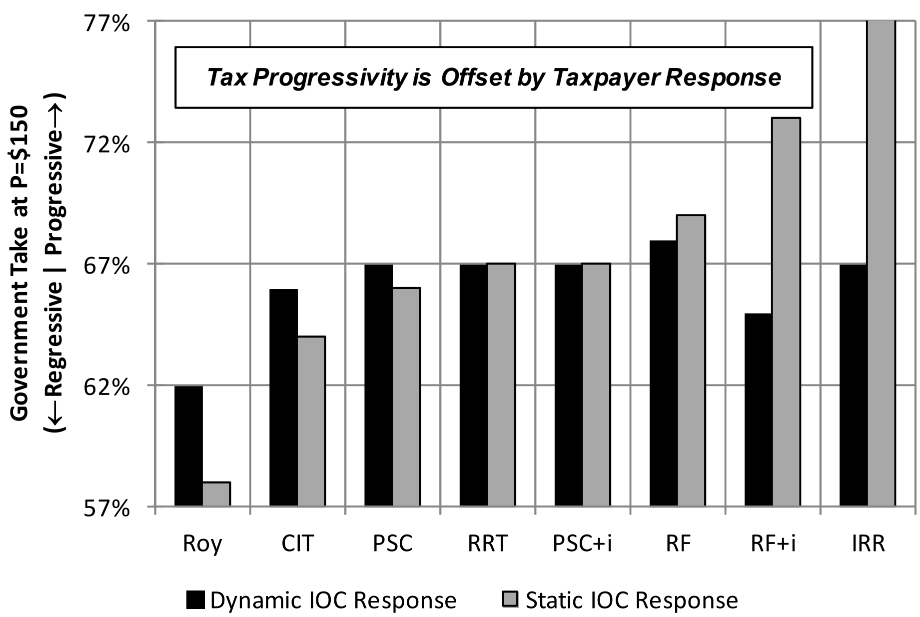

[12] The standard approach in public finance literature is to treat "progressivity" as a static attribute of the tax system with no allowance for behavioral response, but also to recognize that the degree of progressivity will have behavioral effects that may affect realized tax revenues.

According to the static measure of progressivity (light bars), the regimes differ substantially, ranging from the strongly regressive ROY regime to the highly progressive IRR-based PSC. Other (but not all) forms of the PSC, plus the RRT regime, tend to capture a fixed share for the government (equal to about 67 percent). The CIT is only slightly regressive, that being due to rather long depreciation schedules without allowance for cost of carried capital.

But, the picture changes significantly under the second measure of progressivity (dark bars). The taxpayer's response to the different tax instruments, crafted to reduce the burden of each, tends to override the structural progressivity and creates a leveling effect. The ability of the progressive regimes to capture a growing share of rent is largely foiled, as is the tendency of regressive regimes to deliver a smaller share of rent. Government take remains remarkably constant despite the underlying structural differences. Why does this occur? The taxpayer rebalances effort across all the margins of exploitation, avoiding incremental investments that would generate income subject to the progressively higher marginal tax rates in favor of investments that provide offsets to the high marginal rates. In this way, the IOC's behavioral response to higher prices tends to countermand the structural progressivity of these regimes.

The effects of leveling can be seen directly in terms of effective *marginal* rates of tax on the incremental profits that are generated when prices rise from \$100 to \$150, as shown in Figure 10. The IOC's behavioral response reduces the structural 90 percent marginal tax rate of the IRR-based PSC regime to precisely 67 percent. And it increases the structural 45 percent marginal tax rate of the ROY regime to 59 percent. (Intuition: at higher prices, the IOC produces additional barrels at higher cost, and the fixed royalty constitutes a higher share of profits on those high-cost incremental barrels).

Figure 10. Effective Marginal Tax Rates

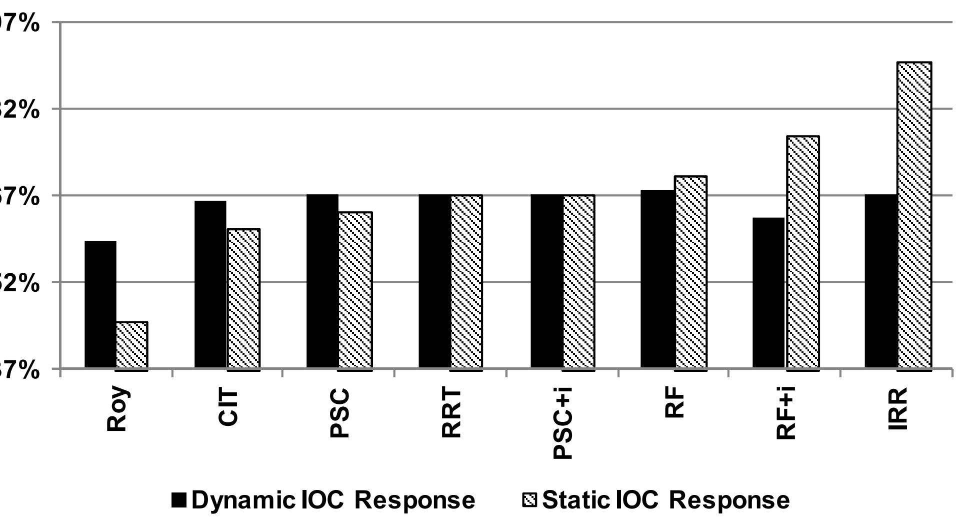

D. Price Volatility and Financial Risk

The structure of the fiscal regime plays an important role in the distribution of risk between HG and IOC. The risk in question is variation in financial returns caused by unpredictable price movements. After the initial investment has been sunk, will the IOC's profit suffer if prices tumble? Or grow if prices rise? And will the government's return be affected likewise? For this purpose, the "static" framework seems most relevant since the IOC would not able to recover sunk investment costs by downsizing its already installed capacity. Moreover, the magnitude of the risk is best measured relative to the size of the expected return since the significance of losing a small absolute amount of profit is greater if expected profits are small to begin with. Therefore, we focus on the coefficient of variation (the standard deviation of return divided by the expected return) as a relevant measure of the risk born by each party.[13] For convenience, we will hereafter use "CV" to refer to this measure of risk.

Any tax regime that captures a fixed share of total profit, independent of price movements, exposes the IOC and HG equally to the consequences of favorable and unfavorable price movements. Each party bears a proportionate share of gains and losses induced by price volatility. Under such regimes, the observed CV of HG and IOC should be equal.

Under a regressive regime, the government levy is relatively fixed and therefore varies less than total profits, which reduces the HG's CV (lower risk for a given reward). In the extreme case of an absolutely fixed levy (e.g., upfront signing bonus), the CV of HG return would be zero and the IOC would be made to bear all price risk.

Under a progressive regime, the roles are reversed: the government relies on instruments designed to capture most of the variation in total profits. Thus, it is the IOC's profit that is stabilized under a progressive regime and this would be reflected in a lower CV for the IOC and a higher CV for the HG.

To measure the risks borne respectively by IOC and HG, we impose a mean-reverting stochastic process to simulate annual oil price movements (replacing the deterministic, constant price path assumed previously), and apply plausible parameter values to capture the range of random variation as well as the speed of reversion to the long-term price trend (from a starting level of $100/barrel, expected annual price change = zero percent with 30 percent volatility, and 30 percent speed of adjustment from any deviation back toward the long-term $100 price level). For each fiscal regime, we simulate 100 independent price paths and

[13] If returns follow a normal distribution, which may not be a bad approximation, then the probability of realizing a NPV within 1 standard deviation of the expected NPV is roughly 68 percent. The reported coefficients of variation therefore indicate (approximately) the width of 68 percent confidence intervals for percentage variations in each party's realized NPV.

calculate the realized NPV of each party, assuming the IOC's initial investment is fixed on the basis of the expected $100 price level.

The results are displayed in Figure 11. *Caveat Emptor*: 100 trials are probably too few to draw sweeping generalizations and these results are very preliminary. It seems important to enlarge this analysis in the future and double-check results.

Figure 11. Risk Sharing (Coefficient of Variation in NPV)

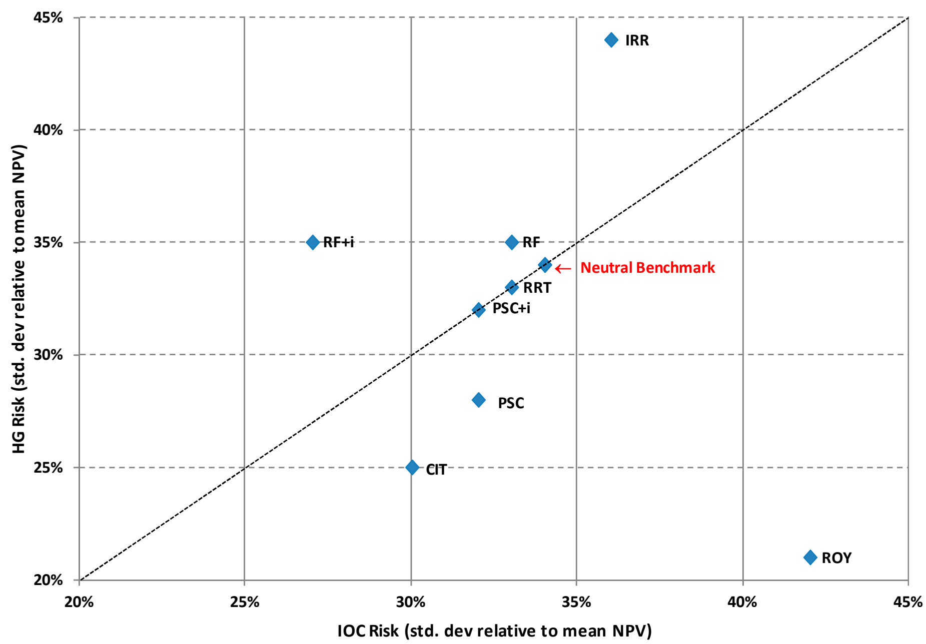

A natural benchmark for judging the distribution of risk is a neutral regime under which HG and IOC each take a fixed share of total profit.[14] With fixed shares, both parties experience the same CV, the level of which does not depend on the fraction by which profits are divided.[15] Because, as shown in the figure, the IOC experiences a 34 percent CV in the NT scenario (where government share is zero percent), it follows that under a neutral tax regime with fixed shares, the standard deviation would be 34 percent of the expected return for each party. Compared to that standard, several of the fiscal regimes considered here allocate less risk to the government. Regimes that subject the HG to greater risk than the IOC are those shown above the dotted diagonal line. As expected, the ROY regime minimizes risk borne by the government and maximizes that borne by the IOC.

[14] This is tantamount to a "Brown tax" in which the government essentially becomes a working-interest partner in the project.

[15] When a random variable is multiplied by a constant, the mean and standard deviation are also multiplied by the same constant, which leaves the CV unchanged.

Only in the case of the IRR-based PSC does the government's share of risk rise appreciably above 34 percent, but that regime has the perverse effect of also raising the amount of risk borne by the IOC. The intuition for this is simple: total returns for each party decline (due to distortions) by more than the standard deviation of those returns (which are less affected by distortions). Other dominance relations among the fiscal regimes are visible in Figure 11, but these reflect something more than just the allocation of risk. We must keep in mind that the underlying investments, and therefore the expected returns, are not being held constant. Rather, each fiscal regime imparts a unique twist to the chosen development program which alters the total amount of risk being allocated between the parties. For this reason, points in the figure do not have the appearance of an efficient frontier.

IOCs also factor risk into their analysis using the so-called profitability index (PI). PI is defined here as the ratio of IOC NPV to the net present value of their maximum cash exposure (NPCE) over the life of a project, which is the amount the IOC stands to lose if the project were aborted at the worst possible time. Thus, PI scales the size of the benefit relative to the amount of money put at risk. Figure 12 shows, for each fiscal regime, PI plotted along the vertical axis and the associated CV plotted along the horizontal axis. Since higher PI is preferred to lower, *and assuming all else equal*, fiscal regimes that register in the upper left corner of the diagram are most favorable to the IOC. Because all fiscal regimes divert cash from the IOC's cash flow stream, they all reduce the profitability index relative to the NT case.

Figure 12. Profitability Index versus Risk

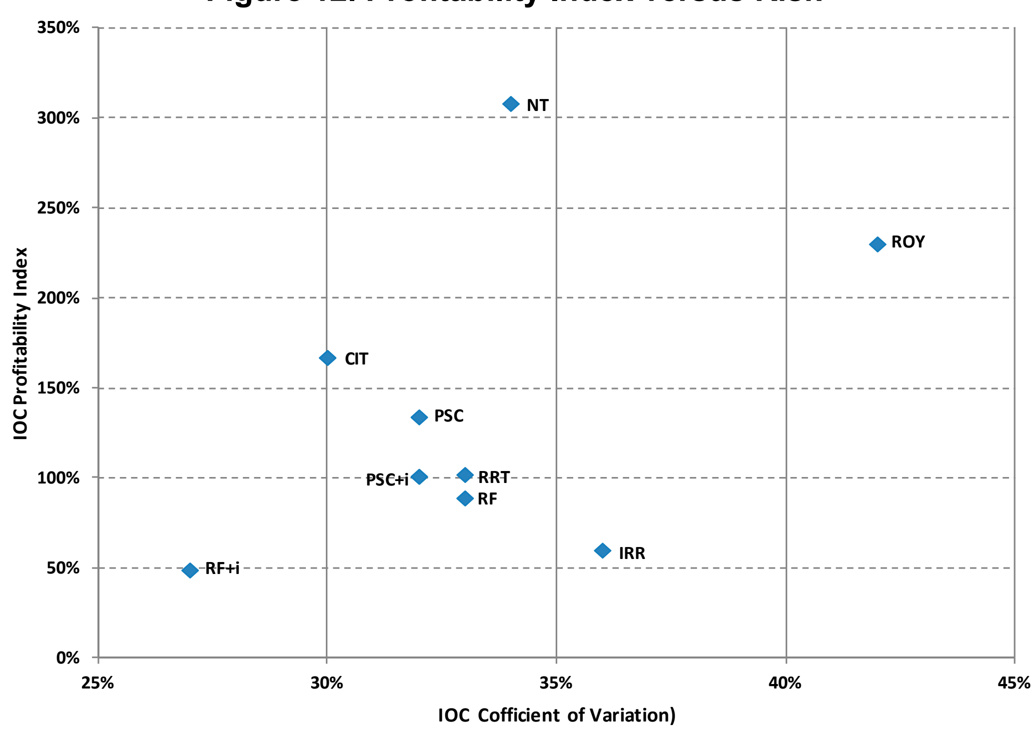

Perhaps surprisingly, the ROY regime achieves the highest PI, but that is only because it dissuades the IOC from making a large initial investment, which directly reduces cash exposure—but that also reduces the opportunity to make a large profit. This illustrates the potential danger of looking at certain investment indices in isolation, since all else is certainly not equal across these regimes. The shortcoming of the ROY regime is clearly apparent in Figure 13, which compares the IOC's simple NPV to CV across regimes. A regime characterized by disincentives for investment is hardly a good thing, even if it generates a high PI, and it is doubtful that any IOC would favor the ROY regime simply on the basis of a high IR.

Figure 13. Net Present Value versus Risk

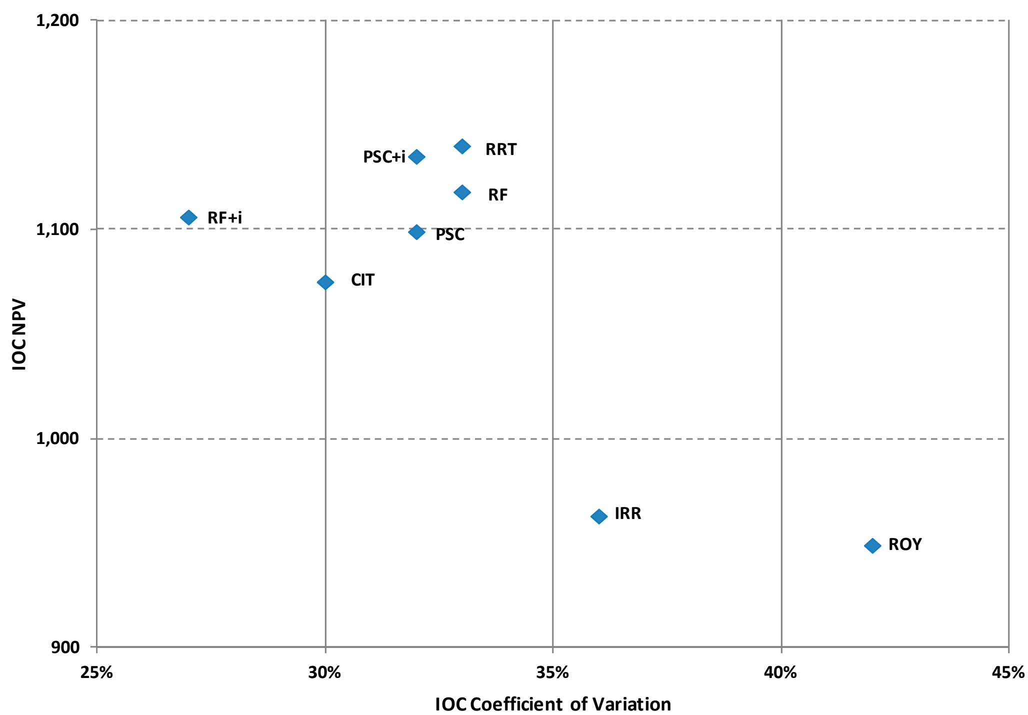

E. Impact of Fiscal Design on the Optionality of Enhanced Oil Recovery

The random price simulations summarized above also shed light on the optionality of the IOC's commitment to enhanced recovery, and the extent to which fiscal design might alter that commitment. Whether the IOC carries out its planned transition to enhanced recovery at the designated time depends on interim price movements, but also on the fiscal burden placed on incremental investments. Although enhanced recovery might have looked profitable at the outset, a prudent investor will rethink the issue as the time draws near. The simulations reported above incorporated this type of optionality by assuming the IOC would (at the time of the proposed EOR implementation) cancel those plans if continuation with primary recovery (which entails no incremental investment) would be more profitable. To this end, we have assumed that, when it comes time to implement EOR, the IOC has perfect foresight regarding future prices and is therefore able to directly compare cash flows with and without the proposed enhanced recovery program. In one sense, this treatment overstates the value of the option since in practice the IOC's decision whether to abort enhanced recovery or not would sometimes be mistaken—and therefore costly. On the other hand, the IOC enjoys other sources of option value that we have not included here (e.g., the IOC might elect to initiate enhanced recovery at an even earlier stage if interim price movements are favorable). A complicated dynamic control model could be applied to precisely value these options, but that is beyond the scope of this analysis. The hope is that the sources of error in our simplified treatment are more or less offsetting, and that the results provide a rough indication of how fiscal regimes interact with the optionality of these investments, and how through this additional channel they might affect the economic efficiency of resource recovery.

Figure 14 summarizes the results of this simulation exercise. Along the horizontal axis, the chart shows the percentage of trials in which initial plans are carried out. Even without tax (NT), unfavorable price movements cause the IOC to abort EOR in 5 percent of the trials. This occurs whenever the economic value of the additional recovered oil would be less than the cost of recovery. From an efficiency perspective, these decisions to forego EOR are justified. Against that standard, however, all the fiscal regimes except RRT artificially raise the probability that EOR will be canceled. The rate of cancellation roughly doubles under the simple PSC and CIT regimes, and more than quadruples (to roughly 20–25 percent) in the case of the ROY and RF PSC. The IRR-based PSC is not represented in this graph since, under that regime, EOR does not appear economically viable even from the beginning.[16] It is worth noting that it is the form of taxation, rather than taxation per se, that reduces the likelihood of EOR. The RRT, which slightly elevates the likelihood of EOR relative to the NT scenario, is a case in point.

The vertical axis of Figure 14 shows the average gain (relative to original IOC NPV) achieved when the option to cancel EOR is exercised. For example, in the NT case, when the option to cancel EOR is exercised, it adds 10 percent on average to the upfront NPV of the field, which is not insignificant. When EOR is cancelled under the other regimes, the average gain ranges from 5 percent (under RRT and CIT) to 19 percent (under RF+i). It must be emphasized that these calculations are based on a relatively small number of Monte Carlo trials in which EOR was in fact cancelled, so the sampling variation may be quite large. Additional simulation studies are needed to draw sharp distinctions between the regimes.

[16] Similar calculations performed for the IRR regime indicate that, although EOR would not have been planned initially, the IOC would in fact choose to implement EOR 36 percent of the time rather than simply abandoning the field at the appointed time.

Figure 14. The Option to Implement Enhanced Oil Recovery

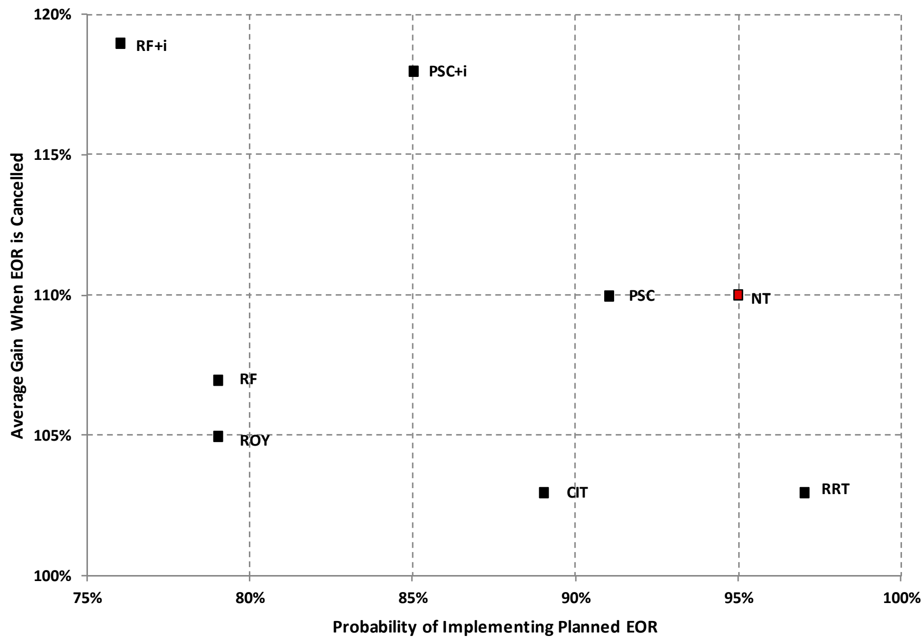

One point is clear, however: the estimated value of the option to cancel EOR is small under the circumstances presumed here. Option value hinges on two factors: the probability of exercise (which ranges from 3 percent to 24 percent across regimes in our trials) and the average gain when exercised (which ranges from 3 percent to 19 percent in our trials). Since both factors are relatively small, so is their combined effect, as shown in Figure 15, which plots the value of the option to cancel EOR as it varies across regimes. Under the NT regime, the option premium amounts to less than 1 percent, primarily because EOR is "deep in the money" without taxes so the option to cancel is rarely used. The same is true of most regimes; given the underlying economic assumptions, EOR is pretty deep in the money. However, to the extent that any particular tax regime erodes the profitability of EOR relative to simple continuation of primary production, the option to cancel EOR gains value.

Figure 15. Value of Option to Cancel Enhanced Oil Recovery

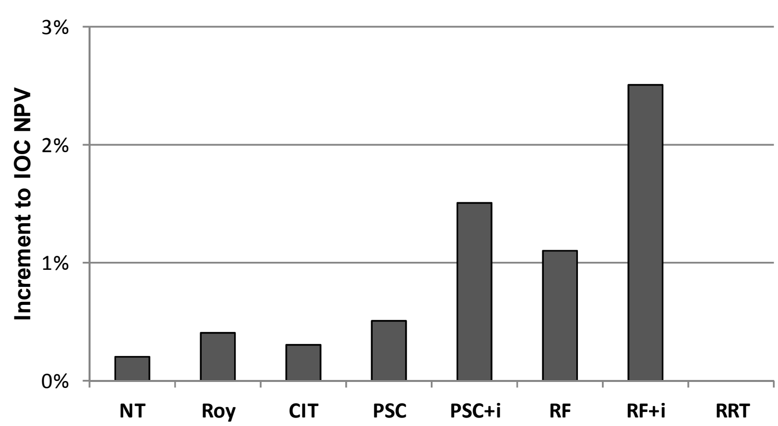

Figure 16 compares IOC NPV computed from the stochastic price simulations with the deterministic IOC NPV presented earlier (Figure 4). On average, the former is 14 percent higher than the latter. However, only a small portion (zero to 2 percent) of that difference consists of option value, as we have just seen. The rest is an artifact of the simulation process itself and stems from the fact that the price of an asset whose return follows a normal distribution tends to rise over time.[17]

Figure 16. Impact of Price Simulations on International Oil Company Net Present Value

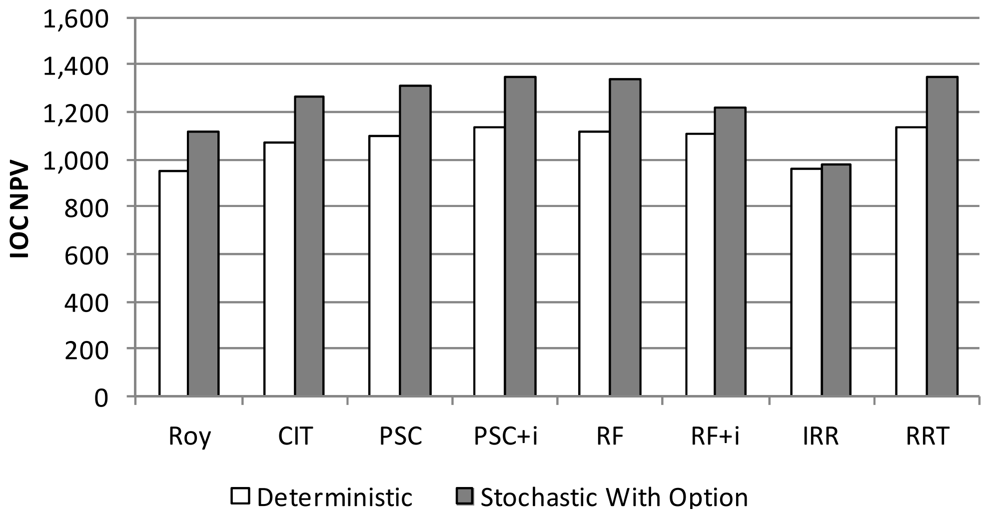

F. Exploration Incentives and Performance

In this section, we turn to results obtained from the integrated exploration and development model. To review that framework, we focus on an IOC that holds the right to drill a series of exploration wells on a given block, subject to a known fiscal regime. The value of each exploratory well depends on the discovery probabilities, which decline in a predictable

---

[17] Our price simulations are based on the assumption that, apart from the mean reversion component, percentage shocks to the price follow a normal distribution with zero mean and volatility = 30 percent: x = ln $(\frac{P_{t+1}}{P_t})\sim N(\mu,\sigma^2)$, with $\mu = 0$ and $\sigma = 0.3$. It follows from statistical theory that the ratio of successive prices must then follow a normal distribution with mean greater than 1, since:

$$y = e^x = \frac{P_{t+1}}{P_t}\sim LN\left(e^{\mu+\frac{1}{2}\sigma^2}, e^{2\mu+\sigma^2} \times \left(e^{\sigma^2}-1\right)\right),$$

where $e^{\mu+\frac{1}{2}\sigma^2}$ = 1.046 given the presumed values of μ and σ. Thus, the price tends to rise in absolute terms over time, which tends to raise returns under the stochastic simulations relative to the deterministic case. The mean reversion component fights against the rising price trend, but cannot completely nullify its effect.

manner after each unsuccessful effort, as well as on provisions of the fiscal regime, and of course on the economic environment (price and cost levels). We assume the IOC will continue with the series of exploratory wells until a discovery is made, or until the expected value of another exploratory well becomes negative—whichever comes first.[18] The value of any discovery anticipated to occur is inferred from the development model discussed previously. Thus, the value of the IOC's right to explore and produce is given by the NPV of the IOC's expected net cash flow stream, beginning with exploration and taking into account those expenses that must be carried forward under the given fiscal regime to the development stage, and incorporating (weighted by discovery probabilities) the ensuing cash flows associated with successful field development, including the potential for EOR. Tax distortions of the scope of exploration are determined endogenously within the model since the IOC's decision whether to abandon the block at any point in the sequence of exploratory wells will depend on how its cash flows (current and future) are taxed.

We begin by looking at the impact of fiscal provisions on the scope of exploration efforts. Figure 17 shows the maximum number of exploratory failures the IOC would tolerate, under each tax regime, before electing to abandon the block. It is assumed that the CIT regime includes a ring fence that prevents the IOC from deducting exploratory costs from revenue generated in other blocks; they may only be carried forward to shelter future revenue on the present block if and when a commercial discovery is made. Under a ring fence, therefore, a decision to abandon the block is also a decision to abandon a tax shelter, and this consideration is factored into the IOC's decision.[19]

Figure 17 reports results from the two alternative assumptions regarding the IOC's information set. Scenario A ("Drill Bit") assumes the IOC places greater weight on the results of exploratory wells than on prior geological evidence. Scenario B ("Geologist") assumes the converse.

[18] An alternative approach would be to assume that multiple fields could exist within the block, and that the IOC would not necessarily cease exploration after the first discovery. An alternative model of discovery probabilities would apply in that case and it would be necessary to incorporate the impact of "sampling without replacement" at the exploration stage.

[19] Without the ring fence, we assume that intangible exploratory drilling costs can (and will) be used to offset revenue from other blocks. We also assume that 80 percent of exploration costs are intangible, which means that even without a ring fence, 20 percent of exploration costs must be carried forward.

Figure 17. Maximum Exploratory Failures before Abandonment

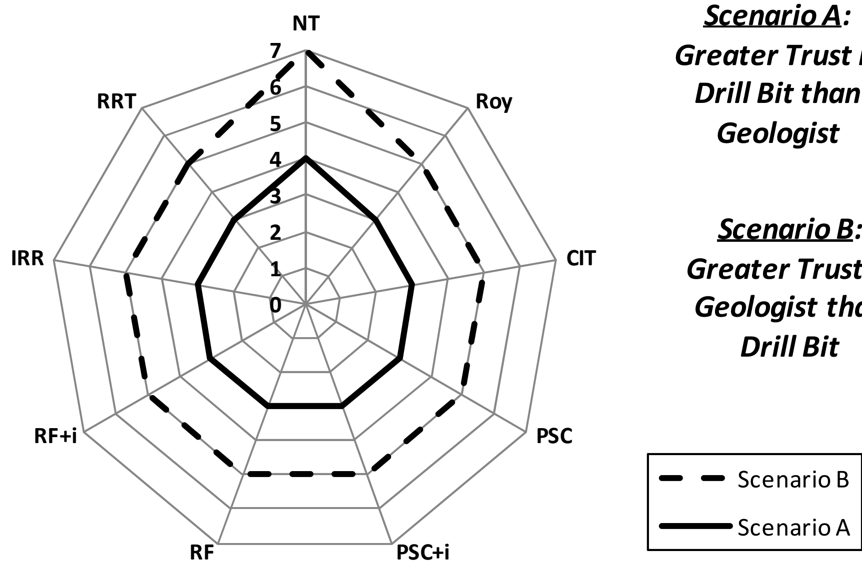

In the case of Scenario A, for example, the figure shows that, although it would take four failures to cause the IOC to abandon the block absent taxes, all fiscal regimes considered here reduce the number of tolerable failures to three, and by cutting short the exploratory effort these fiscal regimes also reduce the chance that hydrocarbons will be discovered. The ring-fence provision is not very onerous under the circumstances assumed here. In the case of Scenario B, all fiscal regimes reduce the maximum number of wells by two; the IOC would drill up to seven exploratory wells in the absence of taxes, but only five if taxes are factored in. (Although not shown here, relaxing the ring fence in the CIT regime would increase the maximum number of wells to six.)

To clarify the influence of the IOC's information set, and how it affects perceived dry hole risk, we review here the distinction between Scenarios A and B. In Scenario A, we assume a technological probability of 70 percent and a geological probability of 50 percent. Together, these determine the chance of success on the first well to be 35 percent (=70 percent x 50 percent), but because more confidence is attached to the drilling outcome, the chance of success drops significantly after each successive dry hole. The situation is different if the IOC places less confidence on drilling outcomes and more on the geological prior. In Scenario B, for example, we assume that the technological probability drops to 50 percent but the geological probability rises to 70 percent. Chance of success on the first well remains 35 percent, as before. But each failed well now has less impact since that "test" of the geology is viewed as less accurate than before, which explains why the IOC would not abandon exploration so soon under Scenario B. Table 5 shows how the probability of success develops differently under the two scenarios.

Table 5. Marginal Chance of Exploratory Success

	Scenario A Emphasis on Drill Bit	Scenario B Emphasis on Geology
Well #1	35.0%	35.0%
Well #2	10.5%	17.5%
Well #3	3.2%	8.7%
Well #4	0.9%	4.4%
Well #5	0.3%	2.2%
Well #6	0.1%	1.1%
Well #7	0%	0.6%
TOTAL	50.1%	69.5%

Another implication of the different belief structures can be drawn from the difference in success probabilities shown in Table 5: since the probabilities fall more slowly under Scenario B, the perceived overall chance of success is higher (69.5 percent vs. 50.1 percent) if the IOC's faith in the geological model is strong enough to dominate early drilling failures. This factor translates directly into a higher IOC NPV—even though the initial dry hole risk is the same, as shown in Figure 18. The figure plots "full-cycle" NPV, i.e. the IOC's combined return from the entire sequence of exploration, development, and production that is expected to occur. Confidence in geology (Scenario B) adds roughly 30–50 percent to full-cycle profit, despite equal initial risk of failure. The impact of the respective fiscal regimes is also apparent: relative to the NT NPV of $3,181 million (in Scenario B), these fiscal regimes reduce the IOC's expected profit by at least half.

Figure 18. Full Cycle International Oil Company Net Present Value

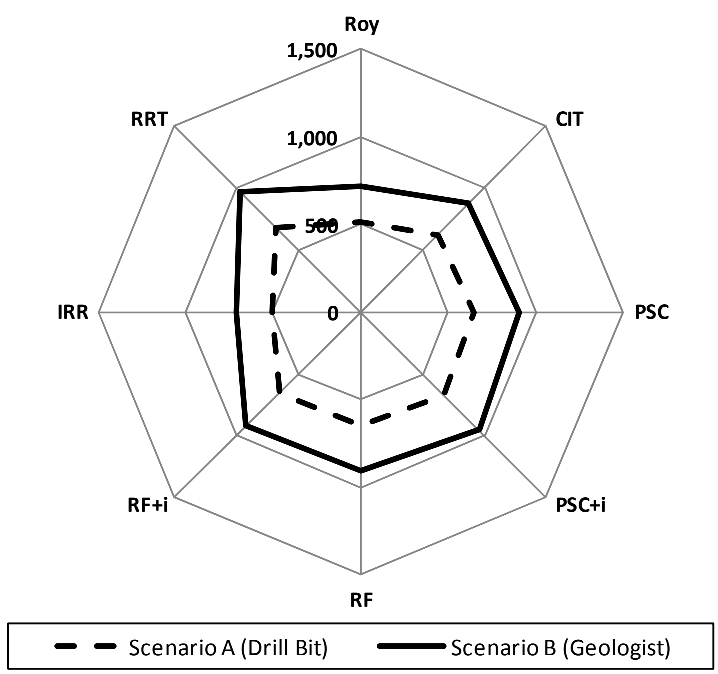

Of course, not all of the IOC's loss is captured by government; some of the potential value of the field is simply forfeited (wasted) due to the distortive impact of fiscal provisions on exploration and development activities. A very simplified indication of that distortion is provided in Figure 19, which summarizes how each regime limits the scope of exploratory drilling (horizontal axis) and, if there were to be a discovery, also limits the resource recovery factor (vertical axis).[20] These are not the only factors related to economic efficiency (timing of recovery is also important), but they are salient indicators of what distortion might look like on the ground.

Figure 19. Distorted Resource Exploration: Exploration and Development Stages

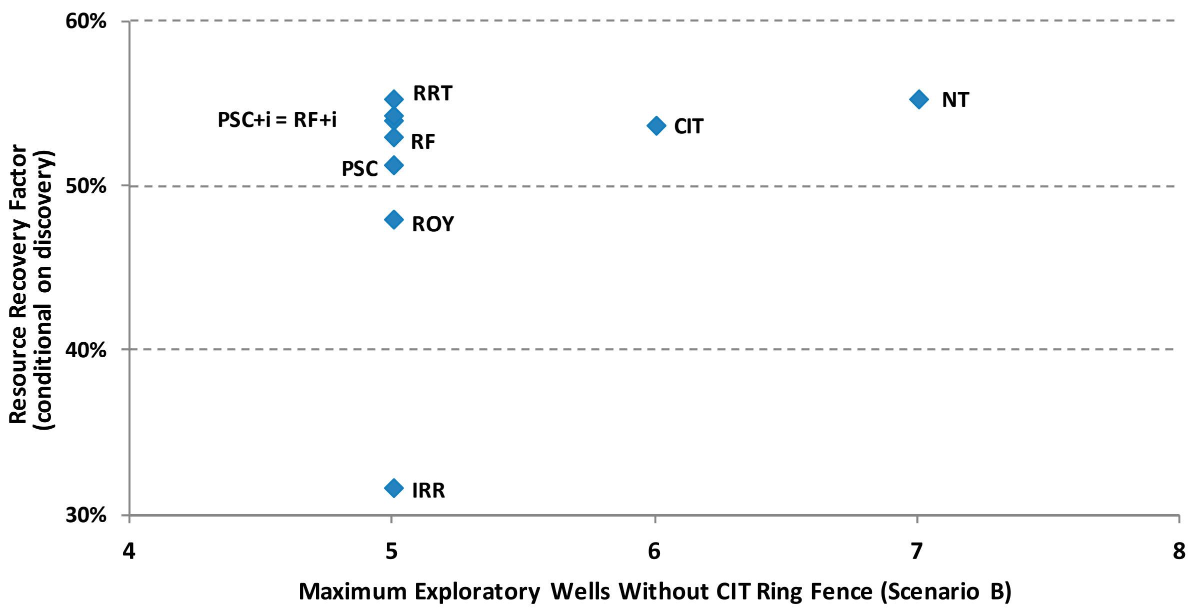

To sum up, Figure 20 shows the combined effect of tax-induced exploration and development distortions on the full cycle value of the enterprise, under the assumptions of Scenario B. Without distortions, the total expected value of the resource play would be $3,181 million (NT regime). The RRT regime comes the closest to that benchmark, at $3,152, thereby achieving 99 percent efficiency. Except for IRR, the other variants on PSCs also come close, ranging from 97 percent to 98.9 percent efficiency. The ROY, CIT, and IRR regimes do not perform as well, at 82 percent, 94 percent, and 83 percent efficiency, respectively.

[20] The figure displays results obtained under Scenario B ("geology") and without a ring fence. Adding the ring fence would reduce the maximum number of wells under the CIT regime from six to five.

Figure 20. Tax Impact on Total Resource Value (Full Cycle)

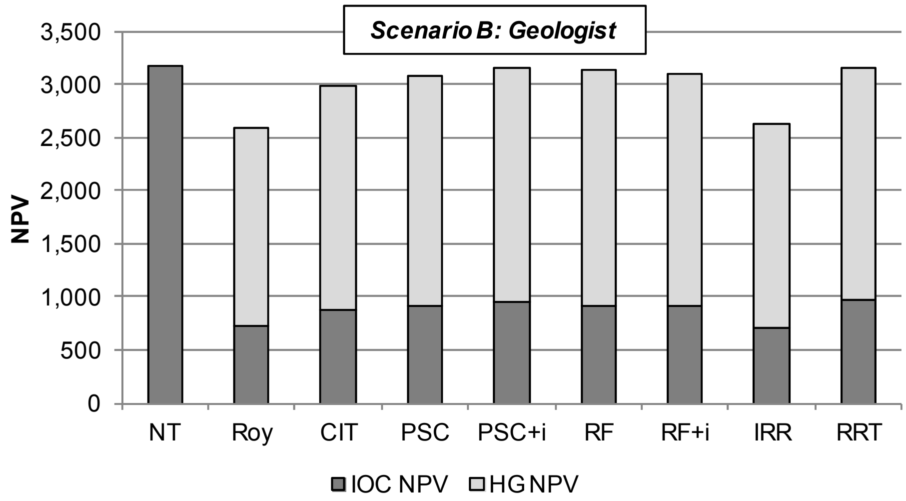

Government take computed in terms of full-cycle net revenues is shown in Figure 21. In every case, the government's share exceeds 67 percent, although each regime was calibrated to capture exactly two-thirds of profits earned at the development stage—post-discovery. The difference is due to the fact that none of these regimes allows full recovery of exploration costs.

Figure 21. Government Take (Full Cycle)

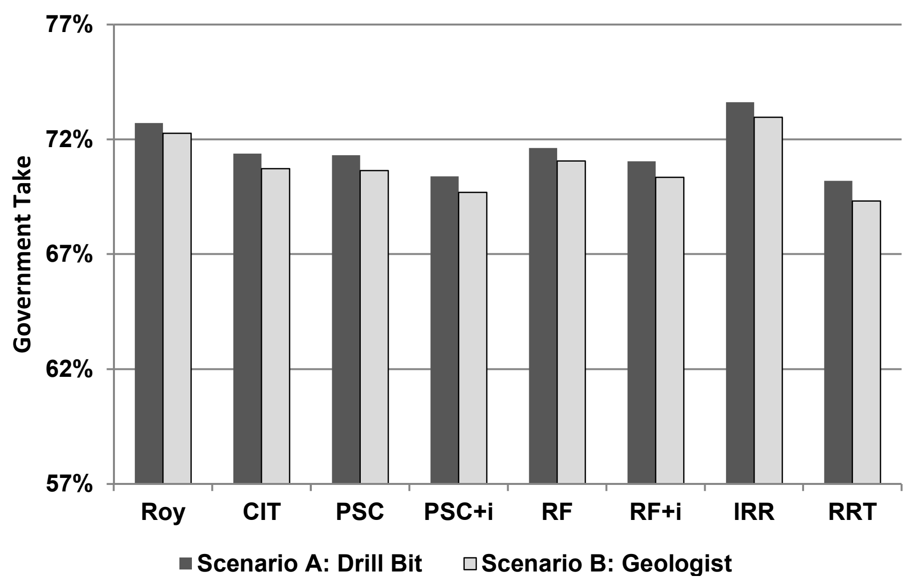

V. Conclusion

This paper has demonstrated how a parsimonious model of petroleum exploration and development can be applied to more fully understand tax distortions and the performance of alternative fiscal regimes. The analysis presented here was mostly illustrative in nature and did not delve very comprehensively into the range of potential applications. For example, we considered only a set of highly profitable petroleum prospects that are to a large degree robust with respect to the level and form of taxation. It should be expected that similar analysis of truly marginal fields would reveal additional insights about the comparative effects of alternative tax instruments. The basic model could easily form the basis of further Monte Carlo experiments to study more systematically the value of optionality in the development program, and how it varies across fiscal regimes. The model could also be used in the customary way to trace fiscal impacts on the thresholds for feasible resource development (e.g., minimum economic price, minimum commercial field size, and maximum development cost).

Although the model comprises a highly simplified abstraction of the actual exploration and development process, it manages to incorporate many factors and some of the key tradeoffs that would influence an investor's investment behavior. It is one of few economic models that explicitly recognize the role of enhanced oil recovery, and since EOR is already a leading source of reserve additions in many parts of the world, it will be useful to study further the potential distortions that may result from one form of taxation or another. It is also one of few economic models that treat the impact of taxation on exploration and development in an integrated manner consistent with an investor's joint optimization of investments at both stages of the process. And because the model is simple, it is also user-friendly, which facilitates application to a broad range of problems without raising computational challenges or prohibitive data requirements.

References

Adelman, M. A., 1990, “Mineral Depletion, with Special Reference to Petroleum,” *The Review of Economics and Statistics*, Vol. 72 (February), pp. 1–10.

Blake, A. J., and M. C. Roberts, 2006, “Comparing Petroleum Fiscal Regimes Under Oil Price Uncertainty,” *Resources Policy*, Vol. 31 (June), pp. 95–105.

Campbell, H. F., and R. K. Lindner, 1985, “A Model of Mineral Exploration and Resource Taxation,” *The Economic Journal*, Vol. 95 (March), pp. 146–60.

Conrad, R. F., and R. B. Hool, 1984, “Intertemporal Extraction of Mineral Resources Under Variable-Rate Taxes,” *Land Economics*, Vol. 60 (November), pp. 319–27.

Daniel, P., B. Goldsworthy, W. Maliszewski, D. Mesa Puyo, and A. Watson, 2010, “Evaluating Fiscal Regimes for Resource Projects: An Example from Oil Development,” in *The Taxation of Petroleum and Minerals: Principles, Problems, and Practice*, ed. by P. Daniel, M. Keen, and C. McPherson (London: Routledge).

Helmi-Oskoui, B., R. Narayanan, T. Glover, and K.S. Lyon, 1992, “Optimal Extraction of Petroleum Resources: An Empirical Approach,” *Resources and Energy*, Vol. 14 (September), pp. 267–85.

Hyde, R., and J. R. Markusen, 1982, “Exploration Versus Extraction Costs as Determinants of Optimal Mineral-Rights Leases,” *Economic Record*, Vol. 58 (September), pp. 224–34.

Jacoby, H. D., and J. L. Smith, 1985, “Effects of Taxes and Price Regulation on Offshore Gas,” *Energy Journal*, Vol. 6 (Special Tax Issue), pp. 279–89.

Johnston, D., 2003, *International Exploration Economics, Risk, and Contract Analysis*, (Pennwell: Tulsa).

Johnston, D., D. Johnston, and T. Rogers, 2008, “International Petroleum Taxation,” Supplement, Independent Petroleum Association of America (Washington, IPAA).

Kemp, A. G., 1987, *Petroleum Rent Collection Around the World* (Halifax: The Institute for Research on Public Policy).

_______, 1992, “Development Risks and Petroleum Fiscal Systems: A Comparative Study of the UK, Norway, Denmark and the Netherlands,” *The Energy Journal*, Vol. 13, pp. 17–39.

_______, 1994, “International Petroleum Taxation in the 1990s,” *The Energy Journal*, Vol. 15 (Special issue), pp. 291–309.

Lund, D., 1992, "Petroleum Taxation Under Uncertainty: Contingent Claims Analysis with an Application to Norway," *Energy Economics*, Vol. 14 (January), pp. 23–31.

Panteghini, P. M., 2005, "Asymmetric Taxation Under Incremental and Sequential Investment," *Journal of Public Economic Theory*, Vol. 7 (December), pp. 761–79.

Petroleum Intelligence Weekly, 2009, "Trickier Fields Mean Saudi Costs are Rising," December 21, Energy Intelligence Group, Inc.

Poterba, J., 2010, "The Challenge of Tax Reform and Expanding the Tax Base," *Economic and Social Review*, Vol. 41, pp. 133–48.

Schiozer, R. F., and S. B. Suslick, 2003, "An Evaluation of a Royalty Relief Model for Mature Fields," AAPG International Conference, Barcelona, Available online at: http://www.searchanddiscovery.com/abstracts/pdf/2003/intl/extend/ndx_81913.pdf.

Smith, J. L., 1995a, "On the Cost of Lost Production from Russian Oil Fields," *The Energy Journal*, Vol. 16, pp. 26–57.

_______, 1995b, "Calculating Investment Potential in South America," *World Oil*, Vol. 216 (June), pp. 117–21.

_______, 1997, "Taxation and Investment in Russian Oil," *Journal of Energy Finance and Development*, Vol. 2, pp. 5–23.

_______, 2005, "Petroleum Property Valuation: The Option to Drill Again," *The Energy Journal*, Vol. 26, pp. 53–68.

_______, 2012, "Issues in Extractive Resource Taxation: Review of Research Methods and Models," IMF Fund Working Paper, forthcoming (Washington: International Monetary Fund).

Smith, J. L., and J. L. Paddock, 1984, "Regional Modeling of Oil Discovery and Production," *Energy Economics*, Vol. 6 (January), pp. 5–13.

Tordo, S., 2007, "Fiscal Systems for Hydrocarbons: Design Issues," World Bank Working Paper No. 123 (Washington: World Bank).

Total Exploration and Production, 2009, "EOR: Maximizing Recovery Factors" in *The Know-How Series* (Paris: Total S.A.) Available online at http://www.total.com/MEDIAS/MEDIAS_INFOS/3404/EN/TOTAL-EORbis-GB2.pdf.

Triest, R. K., 1998, “Econometric Issues in Estimating the Behavioral Response to Taxation: A Nontechnical Introduction,” *National Tax Journal*, Vol. 51 (December), pp.761–72.

Uhler, R. S., 1979, “The Rate of Petroleum Exploration and Extraction,” *Advances in the Economics of Energy and Resources*, Vol. 2, ed. by R. S. Pindyck, (Greenwich: JAI Press, Inc.), pp. 93–118.

Van Meurs, P., 1988, “Financial and Fiscal Arrangements for Petroleum Development: An Economic Analysis,” in *Petroleum Investment Policies in Developing Countries*, ed. by T.W. Wälde and N. Beredjick, (London: Graham and Trotman, Inc.).

________, 2012, “Policy Options for Alaska Oil and Gas,” Presentation to the Alaska Senate Finance Committee, February, Available online at: http://aksenate.org/press/021312_PvM_Presentation.pdf.

Zhang, L., 1997, “Neutrality and Efficiency of Petroleum Revenue Tax: A Theoretical Assessment,” *The Economic Journal*, Vol. 107 (July), pp. 1106–20.