FINANCIAL HEDGING OF OPERATIONAL FLEXIBILITY

MICHAEL LUDKOVSKI

ABSTRACT. We extend the framework of real options to value the compound timing option owned by a manager of an industrial asset. The operator has control over the production modes, but faces operational constraints which introduce path-dependency. Moreover, the operator is only able to imperfectly hedge her income on the futures market. Using an exponential indifference valuation approach we construct a combined stochastic control formulation that merges the problems of optimal switching and indifference pricing in incomplete markets. We then present an iterative scheme for valuing operational flexibility which in particular shows additivity of indifference value over time. After discussing details of numerical implementation, we illustrate our results with several numerical examples and comparative statics.

JEL code: D81, C61, G13

MSC 2000 code: 90B50, 93E20, 91B28, 60G40

Date: First Version: May 2006. This version: October 16, 2007. I thank Patrick Cheridito, Vicky Henderson, Dmitry Kramkov, Sasha Stoikov and seminar participants at University of Florida for useful discussions.

Key words and phrases. operational flexibility, indifference pricing, compound options, stochastic income.

1. INTRODUCTION

Optimizing operation of industrial assets is one of the key problems studied by management science. The manager in charge has control over the operating modes and attempts to maximize expected profit which is tied to a stochastic state variable Y. Thus, she holds a sequence of timing "real options" on the possible revenues. A common example is a commodity production asset with Y_t representing the current commodity price. In this case, the manager has a series of start-up/shut-down options on production. Suppose that the manager can adjust the production regime every Δt time-units and has a choice of I possible regimes with associated cashflows $\psi_i(t, Y_t)$, $i = 0, \dots, I-1$. The classical theory (Dixit and Pindyck 1994, Eydeland and Wolyniec 2003) then implies that the value of the above flexibility on a time horizon $[0, T]$, $T = M\Delta t$ is given by

$$\bar{V}(y) = \mathbb{E}^{\mathbb{Q}}\left[\sum_{m=1}^{M} \max_i \mathrm{e}^{-rm\Delta t}\psi_i(m\Delta t, Y_{m\Delta t}) \cdot \Delta t \,\middle|\, Y_0 = y\right], \tag{1.1}$$

where $\mathbb{Q}$ is the risk-neutral pricing measure for the $\{Y_t\}$-market. Thus, computing the associated value is reduced to pricing a series of chooser Call options on Y. It has been long recognized that (1.1) will in fact *overestimate* the production value due to two crucial phenomena ignored by the classical theory:

(a) The manager faces operational constraints that limit her flexibility;

(b) The asset cannot be perfectly hedged, negating the replication premise underlying risk-neutral valuation. Moreover, the manager is risk-averse and will choose strategies that reduce risk.

Property (a) implies that instead of the sequential Call options of (1.1) one must consider exotic, and particularly *path-dependent* options held by the manager. In fact, for a full account of the constraints, one must work with an entire operational (impulse) control $\xi = (\xi_t)_{0 \le t \le T}$ which represents the dynamic sequence of managerial decisions. After the early seminal paper of Brennan and Schwartz (1985), these issues have been addressed by several papers under the rubric of *optimal switching*. The pde-based approach of quasi-variational inequalities has been considered in Brekke and Øksendal (1998) and studied more thoroughly by Zervos (2003). A probabilistic method was first taken up by Yushkevich (2001) in discrete time; continuous-time versions were then analyzed by Hamadène and Jeanblanc (2007) in the framework of backward stochastic differential equations (BSDEs) and by Carmona and Ludkovski (2005) and Dayanik and Egami (2004) using Snell envelope techniques.

Property (b) arises due to fragmentation of commodity markets as a result of geographical and physical characteristics of the products. Consequently, the manager faces a *basis risk* between the actual commodity produced and the standardized traded contract used for hedging. Thus, the problem must be represented in terms of optimal investment in an incomplete market with (controlled) stochastic income. The risk-preferences of the manager are also important. For instance, poor performance might result in a management re-shuffle and/or worker lay-offs, measures that are undesirable and correspond to a large negative externality. Alternatively, a streak of losses could lead to a credit crisis and rating downgrade of the firm. As a result, the management is likely to place extra emphasis on avoiding bad outcomes and will eschew excessively risky decisions. A natural way of incorporating manager preferences is to apply utility-valuation via an indifference pricing mechanism. The problem of utility maximization with *exogenous* stochastic income has been analyzed in El Karoui and Jeanblanc (1998), Henderson (2005) and with a consumption control in Miao and Wang (2007). An extension beyond expected utility has been recently considered in Klöppel and Schweizer (2007). The related problem of indifference pricing of (European) random endowments has also been extensively studied, see e.g. Henderson (2002), Musiela and Zariphopoulou (2004), Stoikov and Zariphopoulou (2005), Ilhan et al. (2006). The common feature of all these papers is an introduction of a nonlinearity to account for the manager's risk aversion.

1.1. **Combined Formulation.** In this paper, we merge the two aforementioned strands of literature to extend (1.1) in a way that explicitly incorporates operational constraints, imperfect hedging possibilities and risk-preferences of the manager. Thus, we inject risk-preferences into an optimal switching model, or conversely add an *endogenously* controlled stochastic income to the problem of utility maximization. While traditionally the operational and financial arms of the firm have been valued separately, to extract maximum benefits a "holistic" global approach is needed. This is achieved by a combined stochastic control framework which provides a coherent way of analyzing the joint behavior of the manager. The use of a fully dynamic setting properly reflects operational constraints while correctly pricing traded/non-traded risks. Our model is robust and can incorporate many practical extensions; moreover it is computationally tractable, which allows us to give several numerical illustrations. In particular, we use the examples to highlight the interplay between the financial and operational components and to study the role of various constraints.

While this project was completed we became aware of a parallel independent work by Porchet et al. (2007). They consider a very similar financial setting, but frame their model in the context of backward stochastic differential equations. We give a detailed comparison of the two approaches in Remark 4.

1.2. **Case Study.** Before proceeding, let us discuss a case-study that will be used as motivation for further analysis. Consider an oil producer that operates several deep sea oil platforms, extracts oil of specific grade 'Y' and sells it on the market. The company management wishes to maximize risk-adjusted profit on some planning horizon of T years (e.g. for an annual plan). The profit depends on the amount of oil extracted, i.e. the operating policy, as well as on the contemporaneous price of oil, which is random and unpredictable.

To achieve its goals the firm has access to two channels. First, the company controls its own production regime. Thus, when oil prices are lower than extraction costs, the manager has the option to shut down production to minimize losses. Conversely, when commodity prices are high, the company can take full advantage by running all platforms at maximal capacity. This operational flexibility is limited by various *engineering* constraints:

- Changing the production mode is costly: one must dispatch workers to start/stop the platforms and coordinate with the nearby oil pipelines;
- Changing the production mode takes time: the above dispatch takes several days until oil is flowing at the full rated capacity;
- The firm may have market power. While the firm is a negligible player in the global oil market, there are hundreds of specific oil grades and the firm is likely to be a major producer of grade 'Y'. Thus, if grade 'Y' is thinly traded, increasing the firm's production will tend to depress local prices.

Second, the firm has access to the oil futures markets where it can *hedge* its revenues. Financial trading can mitigate the risk associated with uncertain future prices, but also brings risks of its own. In North America the only liquidly traded oil contract is the New York Mercantile Exchange (Nymex) futures based on the West Texas Intermediate (WTI) oil grade. However, less than 10% of all oil produced in US is of WTI grade, and in particular the firm under consideration produces oil of grade 'Y'. While the WTI and 'Y' prices are likely to be highly correlated, hedging is still imperfect and any trading strategy exposes the firm to the residual basis risk.

Given this setting, the manager wishes to find an optimal production policy ξ and an optimal financial trading policy π that would maximize her expected risk-adjusted income over the planning horizon $[0,T]$. Moreover, she is interested in understanding the components of this value, namely the relative costs of constraints, the benefit of each flexibility and the respective synergies in case of potential strategic opportunities.

The resulting model is also applicable to many other economic setting beyond commodity production. Let us mention here management of industrial plants that have fluctuating input costs, labor force administration (with the stochastic factor representing demand), and multi-stage capacity expansion budgeting.

The rest of the paper is organized as follows. Section 2 rigorously constructs the associated control problem; Section 3 then characterizes the structure of the optimal strategy and gives an iterative expression for the indifference value of operational flexibility. Section 4 provides a complete numerical implementation using a simulation approach, which is then used to present two illustrative examples in Section 5. Besides numerical evidence, comparative statics are also analyzed. Finally, Section 6 summarizes our results and discusses further extensions. Most of the technical proofs are delegated to the Appendix.

2. Mathematical Framework

Let Y_t and S_t denote the prices at time t of the local and reference contracts respectively. Thus in the oil company example, Y_t is the price of the produced grade 'Y' oil, while S_t is the price of reference WTI futures. We assume that $\{Y_t\}$ and $\{S_t\}$ are one-dimensional and that $\{S_t\}$ satisfies an Itô stochastic differential equation (SDE) of the form

$$\mathrm{d}S_t = \mu(t) S_t \,\mathrm{d}t + \sigma(t) S_t \,\mathrm{d}W^1_t, \tag{2.1}$$

where W^1 is a standard one-dimensional Brownian motion on a stochastic basis $(\Omega, \mathbb{F}, \mathbb{P})$, and μ, σ are bounded deterministic functions satisfying $\sigma(\cdot) > \epsilon_\sigma > 0$. Thus, $\{S_t\}$ is a time-inhomogeneous geometric Brownian motion. Due to strong seasonality of commodity markets we will explicitly show time-dependence of parameters throughout the paper. Precise dynamics of $\{Y_t\}$ will be specified later on in (2.3).

Besides producing oil priced at Y_t and having access to the $\{S_t\}$-market, the company also maintains a risk-free bank account that earns interest at rate r_t at time t. For clarity of presentation we take $r_t = r$ to be a fixed constant. The extension to deterministic time-varying interest rates is straightforward.

2.1. **Operational Characteristics.** The process $\{Y_t\}$ is used to define the income flow for each operating regime of the asset. We postulate that there are a total of I different operating regimes that we label for convenience as $\mathbb{Z}_I \triangleq \{0, 1, \cdots, I-1\}$. The ordering might indicate the production level (e.g. "offline", "50% capacity", "maximum capacity", etc.), but in general is completely symbolic. For each operating regime $i \in \mathbb{Z}_I$, there is a corresponding (possibly negative) income flow at instantaneous rate $\psi_i(t, Y_t)\mathrm{d}t$. In the case of the oil producer, $\psi_i(t, Y_t)$ represents the nominal value of the oil sold, subject to the assumption that every barrel extracted is immediately sold at prevailing market price. We impose the following

Assumption 1. *For each $i \in \mathbb{Z}_I$, the payoff rate $\psi_i : [0,T] \times \mathbb{R}_+ \to \mathbb{R}$ is uniformly bounded, $|\psi_i(t,y)| \leqslant C_\psi$ for some constant C_ψ. In addition, ψ_i is continuous and has a continuous derivative, $\psi_i \in \mathcal{C}^1_b([0,T] \times \mathbb{R}_+)$.*

Remark 1. Assumption 1 is strictly speaking violated for most practical cases which typically involve income flows ψ_i that are *linear* in the commodity price Y_t. However, one could always bound ψ_i while ensuring that the economic accuracy of the model is not affected, see e.g. a similar Remark 1.2 in Ilhan et al. (2006). Bounded income rates guarantee that operational profits are always finite; this is certainly economically plausible.

An operational strategy ξ is a double sequence $(\tau_k, \xi_k), k = 0, 1, 2, \ldots$, with $\xi_k \in \mathbb{Z}_I$ representing the sequence of chosen production regimes and $0 \triangleq \tau_0 \leq \tau_1 \leq \cdots \leq T$ representing the times of operating regime changes (from now on termed *switching times*). The entire strategy is then the right-continuous $\xi \colon [0, T] \times \Omega \to \mathbb{Z}_I$ with $\xi_t = \xi_k$ if $\tau_k \leqslant t < \tau_{k+1}$ or

$$\xi_t = \sum_{\tau_k < T} \xi_k \cdot \mathbb{1}_{[\tau_k, \tau_{k+1})}(t). \tag{2.2}$$

To match the continuous-time nature of the rest of the model, we have assumed that changes in operating regimes can be carried out at any point in time. A discrete-time version that matches the classical formulation in (1.1) will be discussed in Section 3.5.

Once a strategy ξ is selected, the resulting operation has three effects:

(a) Nominal revenue at rate $\psi_{\xi_t}(t, Y_t)\, \mathrm{d}t$ is generated at time t;
(b) Discrete costs are incurred at times of regime switches. We label these as $C_{i,j} \equiv C_{i,j}(t, Y_t)$, for the expense associated with changing the production from regime i to regime j;
(c) ξ affects the dynamics of the local price $\{Y_t\}$, which follows an SDE of the form

$$\mathrm{d}Y_t = a(Y_t, \xi_t)\, \mathrm{d}t + b(Y_t, \xi_t) \cdot (\rho\, \mathrm{d}W_t^1 + \sqrt{1 - \rho^2}\, \mathrm{d}W_t^2). \tag{2.3}$$

With respect to the switching costs $C_{i,j}$ we make the standing

Assumption 2. *For every $i \in \mathbb{Z}_I$, $C_{i,i} = 0$ and $C_{i,j} \geq 0, \forall j \neq i$. Also, for all $i, j, k \in \mathbb{Z}_I$ $C_{i,k} \leq C_{i,j} + C_{j,k}$ and for any $i, j_1, j_2, \ldots, j_n \in \mathbb{Z}_I$, $C_{i,j_1} + C_{j_1,j_2} + \ldots + C_{j_n,i} > \varepsilon_C > 0$.*

Since $C_{i,j}$'s satisfy the triangle inequality, multiple simultaneous switches are ex ante suboptimal. The last item in the assumption means that switching costs are strictly positive over any "cycle" of decisions. The presence of switching costs implies that the initial regime ξ_0 affects future strategies and introduces path-dependency into the operational optimization problem.

Regarding the dynamics of local prices in (2.3), the driving process W^2 is another one-dimensional Brownian motion, independent of W^1 driving (2.1), and $-1 \leq \rho \leq 1$ is the correlation parameter. Typically ρ is close to $+1$, indicating a high degree of positive dependence between the market futures contract and the local 'Y' commodity. We also postulate that

Assumption 3. *For all $i \in \mathbb{Z}_I$, the coefficients $a(y, i)$ and $b(y, i)$ are bounded and uniformly Lipschitz continuous and the volatility $b(y, i) > \varepsilon_y > 0$ is non-degenerate.*

The effect of ξ on $\{Y_t\}$ models the market power potentially exercised by the manager due to the fundamental laws of supply and demand. For instance, it can represent the condition that when production increases, the supply of the output commodity grows, which in turn tends to drive prices down (or conversely the demand of the input commodity shrinks and prices increase). Due to random price fluctuations, this effect is not deterministic, but is instead incorporated into the price dynamics (deterministic effects can be included directly in ψ_i). When we wish to emphasize this dependence of $\{Y_t\}$ on ξ, we will occasionally write $\{Y_t^\xi\}$, and $\{Y_t^i\}$ for $\xi_t \equiv i \in \mathbb{Z}_I$. We denote by $\mathcal{F}_t = \sigma((Y_s)\colon 0 \leqslant s \leqslant t)$ the filtration generated by $\{Y_t\}$ and by $\widetilde{\mathcal{F}}_t = \sigma((S_s, Y_s)\colon 0 \leqslant s \leqslant t)$ the joint filtration of $\{S_t\}$ and $\{Y_t\}$.

Not all controls are acceptable. First, since ξ should be based on current production conditions, τ_k are required to be $\mathcal{F}$-stopping times. For a $\mathcal{F}$-stopping time σ let $\mathcal{S}(\sigma) = \{\mathcal{F}\text{-stopping time } \tau\colon \sigma \leq \tau \leq T\}$ (with $\mathcal{S} = \mathcal{S}(0)$) be the set of all stopping times between σ and T. Then we need $\tau_{k+1} \in \mathcal{S}(\tau_k)$. Also, the production decisions must be done on the basis of information at switch time, i.e. $\xi_k \in \mathcal{D}(\tau_k, \xi_{k-1})$ where

$$\mathcal{D}(\tau_k, \xi_{k-1}) \triangleq \{d\colon \Omega \to \mathbb{Z}_I, \mathcal{F}_{\tau_k} - \text{measurable}, d(\omega) \neq \xi_{k-1}(\omega)\},$$

denotes the set of all $\mathbb{Z}_I$-valued, $\mathcal{F}_{\tau_k}$-measurable random variables that are a.s. different from ξ_{k-1} (in order to be able to call τ_1 truly a switching time). Second, the manager is not allowed to make "too many" changes. Formally, an acceptable ξ should be finite in the sense that $\mathbb{P}[\tau_k < T \ \forall k] = 0$. Because potential operational profits are finite while switching costs are strictly positive, a strategy with an infinite number of switches is sub-optimal anyway. For any stopping time τ, we designate by

$$\mathcal{U}(\tau, T) = \{\xi\colon \mathcal{F}\text{-adapted}, \mathbb{Z}_I\text{-valued right continuous process of } \mathbb{P}\text{-a.s. finite variation on } [\tau, T]\}, \tag{2.4}$$

the set of all admissible operational strategies between τ and T. Assumption 3 implies that (2.3) has a unique strong solution on $[t, T]$ for each $\xi \in \mathcal{U}(t, T)$ and any initial condition $Y_t = y$.

Remark 2. As mentioned in the introduction, management may be subject to *operational inertia*: after changing into a new regime i one must wait some amount δ_i before being allowed to switch the regime again. For instance, an oil platform has ramp-up/ramp-down periods for the pumps during which no new action is possible. Thus, if $\xi_k = i$ then we may also require $\tau_{k+1} \geq \tau_k + \delta_i$. The delay lengths $\delta_i \geq 0$ are additional constraints and prevent the manager from immediately reversing a decision. To ease on presentation we do not incorporate operational inertia at this stage, but will further explore this issue in Section 5.5.

Summarizing, employing a strategy $\xi \in \mathcal{U}(T_1, T)$ on a time interval $[T_1, T_2]$ yields at date T_2 a nominal cumulative revenue of

$$B_{T_1,T_2}(\xi) \triangleq \int_{T_1}^{T_2} \mathrm{e}^{r(T_2-s)} \psi_{\xi_s}(s, Y_s)\, \mathrm{d}s - \sum_{k\geq 1\colon \tau_{k-1}<T_2} \mathrm{e}^{r(T_2-\tau_k)} C_{\xi_{k-1},\xi_k}. \tag{2.5}$$

For the strategy $\xi_s \equiv i$ always using regime i, we will write $B_{T_1,T_2}(i) \triangleq \int_{T_1}^{T_2} \mathrm{e}^{r(T_2-s)} \psi_i(s, Y_s)\, \mathrm{d}s$. Also for later use we note that for the first switching time τ_1 of a control ξ with $\xi_{T_1} = i$ we have:

$$\begin{aligned} B_{T_1,T_2}(\xi) &= \int_{T_1}^{\tau_1} \mathrm{e}^{r(T_2-s)} \psi_{\xi_s}(s, Y_s)\, \mathrm{d}s - \mathrm{e}^{r(T_2-\tau_1)} C_{i,\xi_1} + \int_{\tau_1}^{T_2} \mathrm{e}^{r(T_2-s)} \psi_{\xi_s}(s, Y_s)\, \mathrm{d}s - \sum_{\substack{k\geq 2\\ \tau_{k-1}<T_2}} \mathrm{e}^{r(T_2-\tau_k)} C_{\xi_{k-1},\xi_k} \\ &= \mathrm{e}^{r(T_2-\tau_1)} (B_{T_1,\tau_1}(i) - C_{i,\xi_1}) + B_{\tau_1,T_2}(\xi). \end{aligned} \tag{2.6}$$

2.2. **Financial Hedging Strategies.** In contrast to operational policies ξ which are essentially discrete, a financial trading strategy π is a continuous control. Let π_t denote the dollar amount of contract S_t held at time t, with the remainder invested in the savings account. Let X_t denote the current wealth in the possession of the firm at time t. Given an initial endowment of x, the outcome of any combined strategy (π, ξ) is a wealth process $(X_t \equiv X_t^{t,x,\pi,\xi})_{t\geq 0}$ which obeys $X_t^{t,x,\pi,\xi} = x$ and

$$\mathrm{d}X_u^{t,x,\pi,\xi} = \pi_u \frac{\mathrm{d}S_u}{S_u} + r_u(X_u^{t,x,\pi,\xi} - \pi_u)\,\mathrm{d}u + \psi_{\xi_u}(u, Y_u)\,\mathrm{d}u - \sum_{k\geq 1} C_{\xi_{k-1},\xi_k} \mathbb{1}_{\tau_k = u}. \tag{2.7}$$

Because $\frac{\mathrm{d}S_t}{S_t} = \mu(t)\,\mathrm{d}t + \sigma(t)\,\mathrm{d}W_t^1$, the traded futures price S_t drops out from (2.7) and is omitted from future analysis.

Let $\mathcal{M}^S \triangleq \{Q \ll \mathbb{P}\colon \mathbb{E}^Q[\ln \frac{\mathrm{d}Q}{\mathrm{d}\mathbb{P}}] < \infty,\, \{\mathrm{e}^{-rt}S_t\}$ is a Q-martingale$\} \neq \emptyset$ be the non-empty set of all martingale measures with finite relative entropy. To exclude arbitrage and make sure that (2.7) is well-defined we require

$$\pi \in \mathcal{A}(t,T) \triangleq \Big\{ (\pi_s \equiv \pi(s, S_s, X_s, Y_s))_{t\leq s\leq T} : \int_t^{\cdot} \pi_s \frac{\mathrm{d}S_s}{S_s} \text{ is a } \mathbb{Q}\text{-supermartingale } \forall \mathbb{Q} \in \mathcal{M}^S \Big\}. \tag{2.8}$$

The motivation for the above choice is to have additivity of the admissible financial hedging policy sets $\mathcal{A}(t,s)$: for any $\widetilde{\mathcal{F}}$-stopping time τ, if $\pi^{(1)}, \pi^{(2)} \in \mathcal{A}(t,T)$ then so is $\mathbb{1}_{t<\tau}\pi_t^{(1)} + \mathbb{1}_{t>\tau}\pi_t^{(2)} \in \mathcal{A}(t,T)$. Observe that for this to occur, $\mathcal{A}(t,T)$ must be independent of initial wealth x and the operational strategy ξ. This is one reason why wealth/trading constraints are difficult to incorporate in our framework and why we do not impose the usual constraint of no-bankruptcy $X^{t,x,\pi,\xi} \geq 0$ (in which case the choice of production strategy ξ would influence admissibility of π).

2.3. **Optimization Problem.** We are finally ready to define the optimization problem. Let

$$U(w) = -\exp(-\gamma w), \qquad \gamma > 0.$$

This is the well-known exponential utility with Constant Absolute Risk Aversion (CARA) parameter $\gamma = \frac{U''(w)}{U'(w)}$. It has been widely used in portfolio optimization literature, see e.g. Carmona and Danilova (2003), Henderson (2005), Ilhan et al. (2006), Miao and Wang (2007), Musiela and Zariphopoulou (2004), Stoikov and Zariphopoulou (2005), Zariphopoulou (2001). Further reasons for choosing exponential utility in our model are discussed in Section 3.3. The manager's control problem is to maximize the expected future utility of terminal wealth $V\colon [0,T]\times\mathbb{R}_+\times\mathbb{R}\times\mathbb{Z}_I \to \mathbb{R}$, over all admissible operating and hedging strategies,

$$V(t,y,x,i) = \sup_{\substack{\pi\in\mathcal{A}(t,T)\\ \xi\in\mathcal{U}(t,T)}} \mathbb{E}_{t,x,y,i}\big[U(X_T^{t,x,\pi,\xi})\big], \tag{2.9}$$

where $\mathbb{E}_{t,x,y,i}[\cdot] \triangleq \mathbb{E}\big[\cdot\big|Y_t = y, X_t = x, \xi_t = i\big]$ denotes conditional expectation under the physical measure $\mathbb{P}$ given the state variables at time t. Problem (2.9) assumes for simplicity that the terminal salvage value is zero; a general residual value of the form $G(Y_T)$ can be easily added. It is well-known that because operational gains are uniformly bounded, for every $\pi \in \mathcal{A}(t,T)$ the family $\{\exp(-\gamma X_\tau^{t,x,\pi,\xi})\}_{\tau\in\mathcal{S}}$ is uniformly integrable and therefore the solution of (2.9) is well-defined and finite. Note that for exponential utility, one may equivalently work with the smaller set of admissible strategies,

$$\mathcal{H}(t,T) = \big\{(\pi,\xi)\colon (X_s^{t,x,\pi,\xi}) \text{ is uniformly bounded from below (in } s,\omega)\big\}.$$

2.4. **Indifference Value of Operational Flexibility.** Let

$$\mathbb{U}^0(t,x) \triangleq \sup_{\pi \in \mathcal{A}(t,T)} \mathbb{E}_t\big[U(X_T^{t,x,\pi})\big], \tag{2.10}$$

denote the value function when the business is not present and the company only engages in financial trading. In this case the evolution of the wealth process is simply

$$X_u^{t,x,\pi} = x + \int_t^u \pi_s \frac{\mathrm{d}S_s}{S_s} + \int_t^u r_s(X_s^{t,x,\pi} - \pi_s)\,\mathrm{d}s, \tag{2.11}$$

and we have a standard Merton problem of portfolio optimization. The value of having control of the business on $[t,T]$ is $p \equiv p_{t,T}(y,x,i)$, where the *indifference value* p satisfies

$$V(t,y,x,i) =: \mathbb{U}^0(t, x+p). \tag{2.12}$$

In other words, $p_{t,T}(y,x,i)$ denotes the initial increase in wealth that balances out relinquishing operational control subject to the given initial conditions. Hence, the agent is indifferent between receiving $p_{t,T}(y,x,i)$ dollars immediately or being granted management privileges until T.

Remark 3. Conversely, assuming the point of view of a firm preparing to start financial hedging, we can also assign an indifference value for having access to new financial markets. Namely, defining

$$\hat{V}(t,y,x,i) \triangleq \sup_{\xi \in \mathcal{U}(t,T)} \mathbb{E}_{t,y,i}\Big[U\big(x\mathrm{e}^{r(T-t)} + B_{t,T}(\xi)\big)\Big] = \mathrm{e}^{-\gamma \mathrm{e}^{r(T-t)}x} \cdot \sup_{\xi \in \mathcal{U}(t,T)} \mathbb{E}_{t,y,i}\Big[-\mathrm{e}^{-\gamma B_{t,T}(\xi)}\Big] \tag{2.13}$$

to be the expected profit from just managerial control, the value of financial hedging is the amount $\hat{p} \equiv \hat{p}_{t,T}(y,x,i)$ that solves $V(t,y,x,i) =: \hat{V}(t,y,x+\hat{p},i)$. We study $\hat{p}$ in Section 3.4.

3. Method of Solution

The double optimization in (2.9) and the presence of four state factors should testify to the complexity of our problem. Note that the model is also non-time-homogeneous. This is a key feature of practical applications, not only because of the finite planning horizon T, but also due to inherent seasonality in commodity markets. Nevertheless, thanks to the special structure it is possible to separate the financial hedging and operational management problems and obtain an efficient solution algorithm.

The key simplification occurs because the manager's preferences are over terminal wealth and intermediate income does not affect availability of trading strategies. As a result the problem can be reduced to one where the entire cumulative revenue $B_{t,T}(\xi)$ of (2.5) is received at T. Indeed, one can re-write (2.7) as

$$\begin{aligned} X_T^{t,x,\pi,\xi} &= \mathrm{e}^{r(T-t)}x + \int_t^T \mathrm{e}^{r(T-u)}\pi_u\left(\frac{\mathrm{d}S_u}{S_u} - r_u \mathrm{d}u\right) + \int_t^T \mathrm{e}^{r(T-u)}\Big(\psi_{\xi_u}(u,Y_u) - \sum_{k\geq 1} C_{\xi_{k-1},\xi_k}\mathbb{1}_{\tau_k=u}\Big)\mathrm{d}u \\ &= \mathrm{e}^{r(T-t)}x + \int_t^T \mathrm{e}^{r(T-u)}\pi_u\left(\mathrm{d}S_u/S_u - r_u\mathrm{d}u\right) + B_{t,T}(\xi) \\ &= X_T^{t,x,\pi} + B_{t,T}(\xi), \end{aligned} \tag{3.1}$$

with the $X^{t,x,\pi}$ from (2.11).

3.1. **Separation Principle.** Fixing ξ and making use of (3.1) we see that (2.9) is related to the problem of utility maximization with random endowment. In particular, we recall the following lemmas regarding optimal investment in the incomplete (S, Y) market due to Tehranchi (2004) and Owen and Žitković (2007). Let

$$M_{t,T} \triangleq \frac{1-\rho^2}{2} \int_t^T \frac{(\mu(s)-r)^2}{\sigma(s)^2}\, \mathrm{d}s. \tag{3.2}$$

The quantity $M_{t,T}$ is related to the Girsanov measure change from $\mathbb{P}$ to the minimal martingale measure $\mathbb{Q}$:

$$\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}} = -\int_0^T \frac{\rho(\mu(s)-r)}{\sigma(s)}\, \mathrm{d}W_s^1 - \int_0^T \frac{\rho^2(\mu(s)-r)^2}{2\sigma(s)^2} \mathrm{d}s.$$

The measure $\mathbb{Q} \in \mathcal{M}^S$ is characterized by the property that it makes $\{\mathrm{e}^{-rt} S_t\}$ into a martingale while unaffecting the law of W^2. It also minimizes the relative entropy with respect to $\mathbb{P}$ among all measures in $\mathcal{M}^S$. Thus, the $\mathbb{Q}$-market price of risk associated to W^1 is the familiar Sharpe ratio $(\mu(t)-r)/\sigma(t)$ and the market price of risk associated to W^2 is zero. It follows that under $\mathbb{Q}$,

$$\begin{cases} \mathrm{d}S_t &= S_t(r\, \mathrm{d}t + \sigma(t)\, \mathrm{d}\tilde{W}_t^1), \\ \mathrm{d}Y_t &= \Big(a(Y_t) - \rho \frac{\mu(t)-r}{\sigma(t)} b(Y_t)\Big)\, \mathrm{d}t + b(Y_t) \cdot (\rho\, \mathrm{d}\tilde{W}_t^1 + \sqrt{1-\rho^2}\, \mathrm{d}\tilde{W}_t^2), \end{cases} \tag{3.3}$$

where $(\tilde{W}^1, \tilde{W}^2)$ is a pair of independent $\mathbb{Q}$-Brownian motions. For typographical convenience we will denote the expectation under the measure $\mathbb{Q}$ as $\bar{\mathbb{E}} \equiv \mathbb{E}^{\mathbb{Q}}$.

Let B be a bounded $\mathcal{F}_T$-measurable random variable and consider the utility maximization problem for random endowment B, $\mathbb{U}(x; B) \triangleq \sup_{\pi \in \mathcal{A}(0,T)} \mathbb{E}\left[-\exp(-\gamma(X_T^{t,x,\pi} + B))\right]$.

Lemma 1. *Tehranchi (2004, Theorem 3.2) Suppose $|\rho| < 1$.*

(a) *Let $p^B = \exp(-\gamma(1-\rho^2)B - M_{0,T})$. The function $\mathbb{U}(\cdot; B)$ is given by*

$$\mathbb{U}(x; B) = -\exp(-\gamma \mathrm{e}^{rT} x) \cdot \bar{\mathbb{E}}[p^B]^{\frac{1}{1-\rho^2}}; \tag{3.4}$$

(b) *An optimal strategy for $\mathbb{U}(x; B)$, $\pi^*(B)$ exists and is equal to*

$$\pi^*(B)_t = \frac{1}{\gamma} \frac{\mu(t)-r}{\sigma(t)^2} \mathrm{e}^{-r(T-t)} + \frac{\rho \mathrm{e}^{-r(T-t)}}{\gamma(1-\rho^2)\sigma(t)} \frac{\beta_t}{\bar{\mathbb{E}}[p^B | \mathcal{F}_t]},$$

where (β_t) is the integrand in the $\mathbb{Q}$-martingale representation of p^B,

$$p^B = \bar{\mathbb{E}}[p^B] + \int_0^T \beta_s (\rho\, \mathrm{d}\tilde{W}_s^1 + \sqrt{1-\rho^2}\, \mathrm{d}\tilde{W}_s^2).$$

Lemma 2. *Owen and Žitković (2007, Theorem 4.1,5.1) Moreover, $\mathbb{U}$ satisfies the following properties*

(a) *Dual optimality: the optimal wealth process $\{X_u^{t,x,\pi^*}\}$ is a $\mathbb{Q}$-martingale;*

(b) *Monotonicity: if $B_1 \le B_2$ $\mathbb{P}$-almost surely then $\mathbb{U}(\cdot; B_1) \le \mathbb{U}(\cdot; B_2)$;*

(c) *Concavity: given $\lambda \in [0,1]$ and two claims B_1, B_2 we have $\mathbb{U}(\cdot; \lambda B_1 + (1-\lambda)B_2) \ge \lambda \mathbb{U}(\cdot; B_1) + (1-\lambda)\mathbb{U}(\cdot; B_2)$;*

(d) *Lebesgue Continuity: if (B_n) are uniformly bounded and $B_n \to B$ $\mathbb{P}$-almost surely, then $\mathbb{U}(\cdot; B_n) \to \mathbb{U}(\cdot; B)$.*

Lemma 1 can be used to obtain a solution to the financial trading aspect of (2.9), allowing us to focus on the operational flexibility component. Indeed, for $\sigma \in \mathcal{S}$, let

$$\alpha_\sigma \triangleq \gamma(1-\rho^2)\mathrm{e}^{r(T-\sigma)}.$$

Proposition 1. *Suppose $|\rho| < 1$. The value function V satisfies*

(3.5)
$$V(t,y,x,i) = \exp\left(-\gamma \mathrm{e}^{r(T-t)}x - \frac{M_{t,T}}{1-\rho^2}\right) \cdot \sup_{\xi \in \mathcal{U}(t,T)} \left\{ \bar{\mathbb{E}}\Big[-\exp(-\alpha_T B_{t,T}(\xi)) \Big| Y_t = y, \xi_t = i\Big]^{\frac{1}{1-\rho^2}} \right\},$$

and the corresponding indifference value solves

$$p_{t,T}(y,i) = \sup_{\xi \in \mathcal{U}(t,T)} \frac{-1}{\alpha_t} \ln \bar{\mathbb{E}}\Big[\exp(-\alpha_T B_{t,T}(\xi)) \Big| Y_t = y, \xi_t = i\Big]. \tag{3.6}$$

The indifference value is independent from initial wealth level $X_t = x$, a pleasing fact that allows valuation without worrying about the current cash position of the firm. The case $|\rho| = 1$ is considered in Section 3.4.

Proof. Fix $\xi \in \mathcal{U}(t,T)$ and denote by

$$V^\xi(t,x,y,i) = \sup_{\pi \in \mathcal{A}(t,T)} \mathbb{E}_{t,x,y,i}\big[U(X_T^{t,x,\pi,\xi})\big] = \sup_{\pi \in \mathcal{A}(t,T)} \mathbb{E}_{t,y,i}\left[-\exp\Big(-\gamma(X_T^{t,x,\pi} + B_{t,T}(\xi))\Big)\right],$$

using (3.1). Then from (2.5) and Assumptions 1-2, $|B_{t,T}(\xi)| \le (T-t)\mathrm{e}^{r(T-t)}C_\psi$ is a $\mathcal{F}_T$-measurable bounded random variable and by Lemma 1,

$$V^\xi(t,x,y,i) = \exp\left(-\gamma \mathrm{e}^{r(T-t)}x - \frac{M_{t,T}}{1-\rho^2}\right) \cdot \bar{\mathbb{E}}\Big[-\exp(-\alpha_T B_{t,T}(\xi)) \Big| Y_t = y, \xi_t = i\Big]^{\frac{1}{1-\rho^2}}.$$

Hence,

$$\begin{aligned} V(t,y,x,i) &= \sup_{\xi \in \mathcal{U}(t,T)} V^\xi(t,y,x,i) \\ &= \exp\left(-\gamma \mathrm{e}^{r(T-t)}x - \frac{M_{t,T}}{1-\rho^2}\right) \cdot \sup_{\xi \in \mathcal{U}(t,T)} \left\{ \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{t,T}(\xi)) \Big| Y_t = y, \xi_t = i\right]^{\frac{1}{1-\rho^2}} \right\}. \end{aligned}$$

If there is no production then $B \equiv 0$, and we obtain $\mathbb{U}^0(t,x) = -\exp(-\gamma \mathrm{e}^{r(T-t)}x - \frac{M_{t,T}}{1-\rho^2})$ in (2.10). Comparing with (2.12), we find

$$-\exp\Big(-\gamma \mathrm{e}^{r(T-t)} p_{t,T}(y,i)\Big) = \sup_{\xi \in \mathcal{U}(t,T)} \left\{ \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{t,T}(\xi)) \Big| Y_t = y, \xi_t = i\right]^{\frac{1}{1-\rho^2}} \right\},$$

which after simplification leads to (3.6). $\square$

3.2. **Dynamic Programming Principle.** Since ξ is an impulse control, one expects that V and p, which respectively satisfy (3.5) and (3.6), satisfy a dynamic programming equation. In this section we rigorously establish this fact which reduces to a series of optimal stopping problems and is then used in Section 4 for numerical computations. The results below resemble existing literature on impulse control, see e.g. El Karoui (1981), Lepeltier and Marchal (1984), Dayanik and Egami (2004). However, because we have a multiplicative- rather than the typical additive-cost structure some adjustments are needed, and we present full proofs for completeness in the Appendix.

To begin solving (3.6) we first consider a sequence of restricted control problems. For any stopping time $\sigma \in \mathcal{S}$, let $\mathcal{U}^k(\sigma, T)$ be the set of all production policies that use at most k switches between σ and T:

$$\mathcal{U}^k(\sigma, T) \triangleq \{\xi = (\sigma, \xi_0, \tau_1, \xi_1, \tau_2, \xi_2, \dots, \tau_k, \xi_k) : \tau_i \in \mathcal{S}(\tau_{i-1}), \xi_i \in \mathcal{D}(\tau_i, \xi_{i-1}),\ i = 1, \dots, k\},$$

and define

$$\phi^0(\sigma, Y_\sigma, i) \triangleq \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\sigma,T}(i)) \middle| \mathcal{F}_\sigma\right], \tag{3.7}$$

$$\phi^k(\sigma, Y_\sigma, i) \triangleq \operatorname*{ess\,sup}_{\xi \in \mathcal{U}^k(\sigma,T), \xi_\sigma = i} \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\sigma,T}(\xi)) \middle| \mathcal{F}_\sigma\right], \quad k = 1, 2, \dots. \tag{3.8}$$

The fact that the left hand sides are a function of Y_σ follows from the (strong) Markov property of $\{Y_t\}$. Note that at this point $\phi^k(\sigma, Y_\sigma, i)$ is defined separately for each σ (a more illuminating but cumbersome notation would be $\phi^k(\sigma; i)(Y_\sigma)$), so a priori it is not at all clear what kind of regularity holds for $\{\phi^k(\cdot, Y_\cdot, i)\}_{\sigma \in \mathcal{S}}$. The next lemma shows that ϕ^k satisfy a recursion formula.

Lemma 3. *For all $k \geq 1$, $\sigma \in \mathcal{S}$, $i \in \mathbb{Z}_I$ we have*

$$\phi^k(\sigma, Y_\sigma, i) = \operatorname*{ess\,sup}_{\tau_1 \in \mathcal{S}(\sigma), \xi_1 \in \mathcal{D}(\tau_1, i)} \bar{\mathbb{E}}\left[\exp\bigl(-\alpha_{\tau_1}(B_{\sigma,\tau_1}(i) - C_{i,\xi_1})\bigr) \cdot \phi^{k-1}(\tau_1, Y_{\tau_1}, \xi_1) \middle| \mathcal{F}_\sigma\right]. \tag{3.9}$$

The following lemma proves the regularity of ϕ^k and implies that (3.9) is a pure optimal stopping problem.

Lemma 4. *For every $k = 0, 1, \dots$ and $i \in \mathbb{Z}_I$, there exists a continuous, bounded $\mathcal{F}$-adapted measurable process $\bar{\phi}^{k,i}$, such that $\phi^k(\sigma, Y_\sigma, i) = \bar{\phi}^{k,i}_\sigma$, the process $\bar{\phi}^{k,i}$ evaluated at stopping time σ.*

Thanks to Lemma 4, we may apply the general theory (El Karoui 1981) of optimal stopping to the Snell envelope of (3.9). This implies that an optimal control exists and gives the following explicit characterization of the solution:

Corollary 1. *The supremum in the optimal stopping problem* (3.9) *for ϕ^k is achieved for*

$$\begin{cases} \tau_* = \inf\{s > \sigma \colon \phi^k(s, Y_s, i) = \max_{j \in \mathbb{Z}_I \setminus \{i\}} (\phi^{k-1}(s, Y_s, j) \cdot \exp(\alpha_s C_{i,j}))\}, \\ \xi_* = \min\{j \in \mathbb{Z}_I \setminus \{i\} \colon \phi^k(\tau_*, Y_{\tau_*}, i) = \phi^{k-1}(\tau_*, Y_{\tau_*}, j) \cdot \exp(\alpha_s C_{i,j})\}. \end{cases} \tag{3.10}$$

To relate the previous developments to our indifference values, we define the operator $\mathcal{G} : \mathcal{C}^\infty([0,T] \times \mathbb{R}_+ \times \mathbb{Z}_I) \hookleftarrow$ by

$$\mathcal{G}w(t, y, i) \triangleq \sup_{\tau_1 \in \mathcal{S}(t), \xi_1 \in \mathcal{D}(\tau_1, i)} \frac{-1}{\alpha_t} \ln \bar{\mathbb{E}}_{t,y,i}\left[\exp\Bigl(-\alpha_{\tau_1}\Bigl[B_{t,\tau_1}(i) - C_{i,\xi_1} + w(\tau_1, Y_{\tau_1}, \xi_1)\Bigr]\Bigr)\right]. \tag{3.11}$$

It is easy to check that $\mathcal{G}$ is an increasing bounded operator, which is moreover continuous in the supremum norm. We now iteratively set $p^0(t, y, i) \triangleq -\frac{1}{\alpha_t} \ln(-\phi^0(t, y, i))$ and $p^{k+1}(t, y, i) \triangleq \mathcal{G}p^k(t, y, i)$, $k = 0, 1, \dots$.

Lemma 5. *We have $\phi^k(t, y, i) = -\exp(-\alpha_t p^k(t, y, i))$.*

Comparing with (3.8) and (3.6) we see that $p^k(t, y, i)$ is therefore equal to the indifference value of production on $[t, T]$ under the constraint that at most k managerial regime switches are possible. Note that $p^k(t, y, i)$ is now defined in terms of the recursive optimal stopping problem (3.11), rather than an impulse control problem of (3.8). In the next section we will take advantage of this fact to compute the unconstrained indifference value $p_{t,T}(y, i)$.

3.3. **Iterative Property of p.** In (3.11) we have constructed a series of optimal stopping problems for the indifference value $p^k(t,y,i)$ of the restricted control problem corresponding to (3.8). We now show that in the limit $k \to \infty$ this produces a solution to the original problem (3.6).

Proposition 2. *The sequence (p^k) is increasing and as $k \to \infty$ converges pointwise to p of* (3.6)*.*

$$\lim_{k\to\infty} p^k(t,y,i) = p_{t,T}(y,i).$$

Moreover, the indifference value $p_{t,T}$ is a fixed point of the operator $\mathcal{G}$: $\mathcal{G}p_{t,T}(y,i) = p_{t,T}(y,i)$, so that $p_{t,T}$ satisfies the dynamic programming equation

$$p_{t,T}(y,i) = \sup_{\substack{\tau\in\mathcal{S}(t),\\ \xi_1\in\mathcal{D}(\tau,i)}} \frac{-1}{\alpha_t} \ln \bar{\mathbb{E}}_{t,y,i}\left[\exp\Big(-\alpha_\tau\Big[\int_t^\tau e^{r(\tau-s)}\psi_i(s,Y_s^i)\,ds - C_{i,\xi_1} + p_{\tau,T}(Y_\tau^i,\xi_1)\Big]\Big)\right]. \tag{3.12}$$

Finally, $p_{t,T}(\cdot,i)$ is locally Lipschitz in y and uniformly bounded on $(t,y)\in[0,T]\times\mathbb{R}_+$.

The unusual property of the effective risk aversion parameter $\alpha_t = \gamma(1-\rho^2)e^{r(T-t)}$ increasing as we move backwards in time is to account for time-value of money. Since all cashflows are stated in *nominal* terms, early on the manager is more risk-averse because opportunity costs are higher initially. For instance, the relative cost of paying $C_{i,j}$ is higher for small t when available capital can be invested profitably for a long time. The above proposition immediately implies

Corollary 2. *An optimal joint strategy (ξ^*,π^*) for $p_{t,T}(y,i)$ exists and is explicitly given by: $\tau_0^* = t, \xi_0^* = i$ and*

$$\begin{cases} \tau_{k+1}^* = \inf\{s > \tau_k^* \colon p_{s,T}(Y_s,\xi_k^*) = \max_{j\in\mathbb{Z}_I\setminus\{\xi_k^*\}} (p_{s,T}(Y_s,j) - C_{\xi_k^*,j})\}, \\ \xi_{k+1}^* = \min\left\{ j\in\mathbb{Z}_I\setminus\{\xi_k^*\} : p_{\tau_{k+1}^*,T}(Y_{\tau_{k+1}^*},\xi_k^*) = p_{\tau_{k+1}^*,T}(Y_{\tau_{k+1}^*},j) - C_{\xi_k^*,j}\right\}, \\ \pi_t^* = -\rho\dfrac{b(Y_t)}{\sigma(t)}\partial_Y p_{t,T}(Y_t,\xi_t^*). \end{cases} \tag{3.13}$$

Proof. The structure of optimal ξ^* follows from the representation (3.12) and the corresponding coupled optimal stopping problem. The optimal financial hedging strategy $\pi^* \triangleq \pi^V - \pi^0$ is defined to be the difference between the optimal π^V of (3.5) and the optimal π^0 of (2.10). The latter is known to be (Henderson 2005) the constant $\frac{1}{\gamma}\frac{\mu(t)-r}{\sigma(t)^2}e^{-r(T-t)}$, and the former is obtained by applying the second item of Lemma 1 to (3.12). Thus, using the notation of Lemma 1,

$$\pi_t^* = \rho e^{-r(T-t)}/(\gamma(1-\rho^2)\sigma(t))\cdot\beta_t/\bar{\mathbb{E}}\big[p^{B_{t,T}(\xi^*)}\big|\,\mathcal{F}_t\big].$$

By Assumption 3, the stochastic flow $y \to Y^y(\omega)$ is a diffeomorphism (Øksendal 1998). Coupled with the differentiability of ψ_i's, this implies that for any fixed $\xi \in \mathcal{U}(t,T)$, the map $y \to \bar{\mathbb{E}}_{t,y,i}[\exp(-\gamma(1-\rho^2)B_{t,T}(\xi))]$ is locally Lipschitz. Finally, applying the Clark-Ocone formula gives the representation of π^* in terms of the (generalized) partial derivative of $p_{t,T}$, see e.g. a similar expression in Henderson (2005). □

Thus, for the first optimal switching time τ, the current indifference value of managerial control is equal to the exponential certainty equivalent of best immediate reward $B_{t,\tau}(i)$ plus the future indifference value of remaining control over $[\tau,T]$. An easy generalization of (3.12) shows that, moreover

$$p_{t,T}(y,i) = \sup_{\xi\in\mathcal{U}(t,T)} \frac{-1}{\alpha_t}\ln\bar{\mathbb{E}}_{t,y,i}\left[\exp\Big(-\alpha_\sigma(B_{t,\sigma}(\xi) + p_{\sigma,T}(Y_\sigma,\xi_\sigma))\Big)\right],$$

for *any* $\sigma \in \mathcal{S}(t)$. Consequently, $p_{t,T}$ is *additive* in time in the sense that there is a functional F such that $p_{t,T} = F(B_{t,\sigma} + p_{\sigma,T})$. A similar observation was made in Carmona and Danilova (2003, Section 4) in the context of indifference pricing of a portfolio of options with different expiration dates. This remarkable property is the major reason for our selection of exponential utility. As shown by Cheridito and Kupper (2006), it is the only such example in the framework of expected utility maximization with allowable negative wealth. Note that additivity is natural and allows the manager to carry out 'local' optimization when choosing her production levels; this also trivially holds in (1.1). To numerically implement this algorithm it remains to evaluate the coupled optimal stopping problems appearing in the construction of $p_{t,T}(y,i)$ in (3.12). This is taken up in Section 4.1 using ideas from American option pricing.

The overall structure of Corollary 2 is reminiscent of problems of optimal investment with discretionary stopping (Karatzas and Wang 2000): the manager should optimally hedge her $\{Y_t\}$-income in the $\{S_t\}$-market and then at her discretion "stop", i.e. switch to a different regime and proceed recursively.

Before continuing, let us briefly summarize the effect of model parameters on the indifference value $p_{t,T}(y,i)$. All these results follow easily from the monotonicity and concavity properties of $\mathbb{U}(x;B)$ in Lemma 2 and from the representation (3.6).

Corollary 3. *The following properties hold for the indifference value $p_{t,T}(y,i)$:*

(a) $p_{t,T}(y,i)$ *is increasing in $\psi_i(t,y)$ for any i (ψ_i measures the operational profitability of the firm).*
(b) $p_{t,T}(y,i)$ *is non-increasing in switching costs $C_{i,j}$.*
(c) $p_{t,T}(y,i)$ *is decreasing in the risk-aversion parameter γ.*
(d) $p_{t,T}(y,i)$ *is increasing in the correlation ρ between the traded and production assets.*

The effect of the volatility $b(y)$ of the production price Y on $p_{t,T}$ is ambiguous. While volatile Y increases opportunities to be in high-value regimes, it also increases day-to-day risk and may lead to more frequent switching costs. Thus, the overall impact of $b(y)$ depends on the ensemble of ψ_i's and $C_{i,j}$'s, as well as on the parameters of the (S,Y)-dynamics.

Remark 4. The concurrent work carried out by Porchet et al. (2007) also studies (2.9). However, in that paper the authors characterize V and p as solutions of a system of reflected quadratic BSDEs. In particular, this allows to introduce a certain class of trading constraints for the manager and consider a multi-dimensional setting. The resulting model requires more delicate handling and in Porchet et al. (2007) the main thrust of the paper is therefore proving the existence and uniqueness results related to $p_{t,T}$ and V using analytic BSDE tools. This contrasts with the direct probabilistic method we employ which effectively sidesteps the problem of optimal investment. Because of the different methods, the technical assumptions on state variables and admissible controls made here and in Porchet et al. (2007) are also slightly different.

While we establish the dynamic programming principle for $p_{t,T}$ directly in (3.12), Porchet et al. (2007) only prove the corresponding result for the value function V. Because of their more general model, they cannot explicitly separate the trading and managerial controls and their version of the Bellman equation does not directly translate into an easily implementable numerical method. As a result, compared to this study, Porchet et al. (2007) devote much less space to illustrations and comparative statics.

3.4. **Limiting Cases.** To better understand the mechanics of (3.6) we consider two limiting cases. The first limiting case is $\rho = 1$, which corresponds to a complete market. When $\rho = 1$, the asset 'Y'-cashflows are perfectly correlated with the market benchmark $\{S_t\}$. Consequently, the cumulative asset profit $B_{t,T}(\xi)$ can be replicated using a trading strategy in $\{S_t\}$. Thus, intuition suggests that risk-neutral pricing should be applied. Indeed, in that case, as shown by Henderson (2005), $\mathbb{U}(x;B) = -\exp\Big(-\gamma e^{rT}(x+\bar{\mathbb{E}}[B]) - \int_t^T \frac{(\mu(s)-r)^2}{2\sigma(s)^2}\,\mathrm{d}s\Big)$. Therefore $p_{t,T}(y,i) = \sup_{\xi\in\mathcal{U}(t,T)} \bar{\mathbb{E}}_{t,y,i}\left[B_{t,T}(\xi)\right]$ and the recursive version is

$$p_{t,T}(y,i) = \sup_{\tau\in\mathcal{S}(t),\xi_1\in\mathcal{D}(\tau_1,i)} \bar{\mathbb{E}}_{t,y,i}\Big[e^{r(t-\tau)}\Big\{\int_t^\tau e^{r(\tau-s)}\psi_i(s,Y_s)\,\mathrm{d}s - C_{i,\xi_1} + p_{\tau,T}(Y_\tau,\xi_1)\Big\}\Big]. \tag{3.14}$$

This is precisely the continuous-time version of the original (1.1) adjusted to take into account operational constraints. Such problems have been solved in the optimal switching literature, see e.g. Zervos (2003) and Carmona and Ludkovski (2005). Compared with (3.12), (3.14) can be seen as a linearization that arises in the absence of market incompleteness. Conversely, the introduction of a nonlinear transformation in (3.12) is to account for imperfect hedging opportunities and corresponds to the concept of *nonlinear conditional expectations* in Musiela and Zariphopoulou (2004). This nonlinearity adjusts the expected future profits in order to separate the hedgeable and non-hedgeable components.

At the other extreme, when $\rho = 0$, the local asset price $\{Y_t\}$ evolves independently of the market benchmark $\{S_t\}$. Intuitively that should make financial hedging impossible. Indeed, plugging-in $\rho = 0$ into (3.13) we get $\pi_t^* = 0$. Thus, since it is impossible to hedge operations, the manager does not make any extra investments in the $\{S_t\}$-market in the presence of the asset.

The case $\rho = 0$ is also related to the value of financial hedging mentioned in Remark 3. Indeed, solving for the gain from financial hedging $\hat{p}_{t,T}(y,i)$ and observing that the supremum in (2.13) is the same as that in (3.5) when $\rho = 0$, we find

$$\hat{p}_{t,T}(y,i) = M_{t,T}/\alpha_t + \Big(p_{t,T}(y,i) - p_{t,T}^{(\rho=0)}(y,i)\Big).$$

Thus, the value of access to the $\{S_t\}$-market decomposes into (a) expected direct gains from trading in the reference $\{S_t\}$-contract and (b) increased value of operational income thanks to reduced risk.

3.5. **Discrete Time Formulation.** To be able to compare the value obtained from (3.12) to the original (1.1), one needs to consider a discrete formulation where the manager is required to make switches at pre-specified times belonging to $\mathcal{S}^\Delta \triangleq \{0, \Delta t, 2\Delta t, \cdots, T = M\Delta t\}$. This means that operating mode decisions are made every Δt time units, for example once a day (so that $\Delta t = \frac{1}{365}$), analogous to the distinction between American and Bermudan exercise rights for a vanilla option. The discrete-time version will also be used in the numerical implementation of Section 4.

Let $\mathcal{U}(t,T) \supseteq \mathcal{U}^\Delta(t,T) = \{\xi\colon \tau_k \in \mathcal{S}^\Delta\}$ denote the corresponding set of discretized operating strategies. Now if $\xi \in \mathcal{U}^\Delta(t,T)$, no operational control is possible between t and $t+\Delta t$ so the optimal stopping problems for p reduce to a sequence of one-period decisions. Formally, let $p^\Delta_{t,T}(y,i) = \sup_{\xi\in\mathcal{U}^\Delta(t,T)} \frac{-1}{\alpha_t}\ln\bar{\mathbb{E}}_{t,y,i}\left[\exp(-\alpha_T B_{t,T}(\xi))\right]$ denote the value *just before* time t. Then a production regime decision may be made *immediately* at t or one must wait until (just before) $t+\Delta t$, so that by analogy with (3.12) we obtain

$$p^\Delta_{t,T}(y,i) = \max_{j\in\mathbb{Z}_I} \frac{-1}{\alpha_t}\ln\bar{\mathbb{E}}_{t,y,j}\Big[\exp\big(-\alpha_{t+\Delta t}[p^\Delta_{t+\Delta t,T}(Y_{t+\Delta t},j) + B_{t,t+\Delta t}(j) - e^{r\Delta t}C_{i,j}]\big)\Big]. \tag{3.15}$$

Note that (3.15) has a deterministic optimization over all the possible regime selections since the corresponding decision is made "today". This was the rationale for looking at values just before switching opportunities.

Remark 5. The effect of Δt on the indifference value $p^{\Delta}_{t,T}$ is small. For regular optimal stopping problems Dupuis and Wang (2005) showed that the corresponding error is of order $\mathcal{O}(\sqrt{\Delta t})$ and same bound was obtained in a linear optimal switching problem in Carmona and Ludkovski (2005). Numerically in the examples below we found the effect to be negligible, being less than 1% between $\Delta t = T/100$ and $\Delta t = T/800$.

4. Numerical Implementation

4.1. **Regression Monte Carlo Method.** Since we allow general Y-dynamics in (2.3) and general cashflow rates $\psi_i(t,y)$ a closed-form solution to (3.6) and (3.12) cannot be expected. Therefore we must resort to numerical methods. To do numeric computations, time must be discretized, so that we continue to work with the discrete set $\mathcal{S}^{\Delta} = \{m\Delta t\}_{m=0}^{M}$ and the corresponding set of operating strategies $\mathcal{U}^{\Delta}$. Let $t_1 = m\Delta t, t_2 = (m+1)\Delta t$ be two generic consecutive time steps. As shown in the previous section, finding the indifference value of the production asset $p^{\Delta}_{t_1,T}(y,i)$ hinges on iteratively computing (3.15), which is a distorted conditional expectation of future $p^{\Delta}_{t_2,T}(\cdot,\cdot)$.

To compute (3.15) we shall use a finite-dimensional projection. Let

$$E_{t_1}[i,j](y) \triangleq \bar{\mathbb{E}}\left[\exp\big(-\alpha_{t_2}(p^{\Delta}_{t_2,T}(Y_{t_2},j) + B_{t_1,t_2}(j) - \mathrm{e}^{r\Delta t}C_{i,j})\big)\Big| Y_{t_1} = y, \xi_{t_1} = j\right]. \tag{4.1}$$

We shall approximate this conditional expectation with a projection operator $\hat{E}_{t_1}[i,j](y) \simeq E_{t_1}[i,j](y)$ defined by

$$\hat{E}_{t_1}[i,j](y) \triangleq \sum_{\ell=1}^{N^B} \alpha^*_\ell B_\ell(y), \tag{4.2}$$

with $B_\ell(y)$ being the N^B basis functions in $L^2(\mathbb{R}_+, \mathcal{F}_{t_1})$ and α^*_ℓ the $\mathbb{R}$-valued coefficients chosen to minimize the squared projection error:

$$\alpha^* = \underset{(\alpha'_1,\dots,\alpha'_{N^B})}{\arg\min} \Big\| \sum_{\ell=1}^{N^B} \alpha'_\ell B_\ell - E_{t_1}[i,j] \Big\|^2_{L^2(\mathbb{R}_+)}.$$

The canonical choice is to take $B_\ell \equiv \tilde{B}_\ell(\mathcal{F}_{t_1})$, $\ell = 1,\dots,N^B$ where $\{\tilde{B}_\ell\}_{\ell=1}^{\infty}$ is a complete orthonormal family in $L^2(\mathbb{R}_+)$, e.g. the Hermite polynomials (Longstaff and Schwartz 2001). Empirically, choice of $\{B_\ell\}$ greatly affects algorithm variance and customizing the basis functions to resemble the expected shape of the function $p_{t,T}(\cdot,i)$ is desirable. Because we regress expressions of the form $\exp(-\alpha_t\tilde{g}(Y_t))$, we pick basis functions of the same form $B_\ell(y) = \exp(-\alpha_t g_\ell(y))$, where $g_\ell(y)$ are polynomials. Having five or six basis functions, $N^B = 6$, normally suffices, and having more bases can often lead to worse numerical results due to overfitting. The exponential transformation inside (4.1) causes some numerical instability, especially as it may produce very small values where round-off errors become a concern.

The projection operator $\hat{E}_{t_1}[i,j]$ can be in turn approximated with an *empirical regression* based on a Monte Carlo simulation. This replaces the optimal α with sample $\tilde{\alpha}$. The Monte-Carlo simulation begins by generating for each $i \in \mathbb{Z}_I$ N^p sample paths $\{y^{n,i}_{m\Delta t}\}$ of the $\{Y^i_t\}$ process (recall that this is the 'Y'-price dynamics conditional on the control $\xi_t \equiv i$), with a fixed initial condition $y^{n,i}_0 = y = Y_0$. If ξ does not influence $\{Y_t\}$, then we can use the same set of paths $\{y^n_{m\Delta t}\}$ to compute all the conditional expectations.

We will approximate the overall indifference value by the empirical average $p_{0,T}(y,i) \simeq \frac{1}{N^p}\sum_n p(0,y^{n,i}_0,i)$. Denote the one-period gains by $\tilde{B}_{m\Delta t}(y^{n,i}_{m\Delta t},i)$ which are the sample realizations of (approximation to)

$\int_{m\Delta t}^{m\Delta t+\Delta t} \exp(r(t+\Delta t - s)) \cdot \psi_i(s, Y_s^i)\,\mathrm{d}s$ conditional on starting value $Y_{m\Delta t} = y^{n,i}_{m\Delta t}$. The pathwise values $p(m\Delta t, y^{n,i}_{m\Delta t}, i)$ are computed recursively in a backward fashion, starting with $p(T, y^{n,i}_T, i) = 0$. At step t_1, we regress the time-t_2 pathwise values corresponding to $E_{t_1}[i,j]$ onto the basis functions $B_\ell(y^{n,j}_{t_1})$, i.e. we find $\tilde{\alpha}_\ell(t_1), \ell = 1, \dots, N^B$ minimizing

(4.3)

$$\tilde{\alpha}(t_1) = \arg\min_{\alpha \in \mathbb{R}^{N^B}} \sum_{n=1}^{N^p} \left(\sum_{\ell=1}^{N^B} \alpha_\ell \cdot B_\ell(y^{n,j}_{t_1}) - \exp\left(-\alpha_{t_2}(p(t_2, y^{n,j}_{t_2}, j) + \tilde{B}_{t_1}(y^{n,j}_{t_1}, j) - \mathrm{e}^{r\Delta t} C_{i,j}) \right) \right)^2.$$

After determining $\tilde{\alpha}(t_1)$, this yields a prediction $\hat{E}_{t_1}\big[i,j\big](y^{n,i}_{t_1})$ for the "continuation value" along the (n,i)-th path if an immediate switch to production regime j is taken. Note that the regression uses the paths from the j-regime, but the conditional expectation is evaluated on the paths in the i-regime. The optimal operational decision is made by identifying the index that maximizes $\hat{E}_{t_1}[i,\cdot]$:

$$\xi^{n,i}_{t_1} = \arg\max_j \hat{E}_{t_1}\big[i,j\big](y^{n,i}_{t_1}), \tag{4.4}$$

so that the set $\{n \colon \xi^{n,i}_{m\Delta t} \neq i\}$ identifies all the paths where switching from regime i at $t = m\Delta t$ is beneficial. The overall recursive *pathwise* construction for p is therefore

$$p(m\Delta t, y^{n,i}_{m\Delta t}, i) = \begin{cases} -\frac{1}{\alpha_{m\Delta t}} \ln \hat{E}_{m\Delta t}[i,i](y^{n,i}_{m\Delta t}) & \text{no switch;} \\ -\frac{1}{\alpha_{m\Delta t}} \ln \hat{E}_{m\Delta t}[i,j](y^{n,i}_{m\Delta t}) & \text{switch to } j. \end{cases} \tag{4.5}$$

We call the above scheme Regression Monte Carlo. It was first proposed by Tsitsiklis and Van Roy (2001) in the context of American option pricing. Note that the basic projection method does not guarantee that $\hat{E}_{t_1}$ will be positive, which is a problem given the reverse log-transformation in (4.5). To overcome this issue, one can use a constrained projection or the following more robust method.

4.2. **Simulating Optimal Realized Gains.** The numerical algorithm of Section 4.1 can be improved by exploiting a device first mentioned by Longstaff and Schwartz (2001). Instead of keeping track of the pathwise $p_{t,T}$ we keep track of pathwise *realized gains*. Let $b^{n,i}_{m\Delta t}$ be sample pathwise realizations of $\mathrm{e}^{-r(T-m\Delta t)} B_{m\Delta t,T}(\xi^*)$. Since $p_{t,T}(y,i) = -1/\alpha_t \ln \bar{\mathbb{E}}_{t,y,i}[\exp(-\alpha_T B_{t,T}(\xi^*))]$, $b^{n,i}_{m\Delta t}$ are proxies for $p(m\Delta t, y^{n,i}_{m\Delta t}, i)$ in the previous section. However, applying (2.6) to the discrete-time version of Section 3.5, we obtain as the analogue of (3.15)

$$\mathrm{e}^{-r(T-t_1)} B_{t_1,T}(\xi^*) = \mathrm{e}^{-r\Delta t} B_{t_1,t_2}(\xi^*_{t_1}) - C_{\xi^*_{t_1}, \xi^*_{t_2}} + \mathrm{e}^{-r\Delta t} \mathrm{e}^{-r(T-t_2)} B_{t+\Delta t,T}(\xi^*),$$

which implies the simpler update rule

$$b^{n,i}_{m\Delta t} = \begin{cases} \mathrm{e}^{-r\Delta t}(b^{n,i}_{(m+1)\Delta t} + \tilde{B}_{m\Delta t}(y^{n,i}_{m\Delta t}, i)) & \text{no switch;} \\ \mathrm{e}^{-r\Delta t}(b^{n,j}_{(m+1)\Delta t} + \tilde{B}_{m\Delta t}(y^{n,i}_{m\Delta t}, j)) - C_{i,j} & \text{switch to } j. \end{cases} \tag{4.6}$$

The switching decision is made in direct analogue to (4.1) by replacing $p(\cdot, y^{n,i}_{m\Delta t})$ with $b^{n,i}_{m\Delta t}$ in (4.3) and using the corresponding version of $\xi^{n,i}_{m\Delta t}$ of (4.4). We start with $bn, i_T = 0$ and after backward recursion report at time $t = 0$ $p_{0,T}(y,i) \simeq -\frac{1}{\alpha_0} \ln\big(\frac{1}{N^p} \sum_n \exp(-\alpha_T b^{n,i}_0)\big)$.

Accordingly, we use the conditional expectations solely to decide whether a switch is optimal or not, and propagate back the pathwise profits $b^{n,i}_t$ based on these decision. Consequently, the projection error only has effect to the extent that it implies a wrong switching decision; as long as the operational policy is correct, the pathwise realized gain is computed exactly. This device avoids the non-negativity constraint on $\hat{E}_{t_1}$ and eliminates intermediate projection errors. It also highlights the fact that the performance of the simulation algorithm depends on the accuracy of the constructed approximately optimal operational policy.

4.3. **Error Analysis and Alternative Methods.** Before proceeding to numerical examples, let us make a few remarks regarding algorithm performance. The numerical analysis of the algorithm is complicated due to several layers of necessary approximations. In particular, because the empirical projections $\hat{E}_{m\Delta t}[i,j]$ use the same set of Monte Carlo paths for different $m\Delta t$'s, the resulting errors are correlated. Moreover, while the approximated projection $\hat{E}_{m\Delta t}$ is a global operator depending on the cross section of all the pathwise values at $y^{n,i}_{m\Delta t}$, the Monte Carlo sampling error produces for each n an individual local error in $p(m\Delta t, y^{n,i}_{m\Delta t}, i)$ (or $b^{n,i}_{m\Delta t}$). This error is propagated back in a nonlinear fashion as it causes fluctuations in $\tilde{\alpha}$ for earlier (in time) regressions.

A complete error analysis for our algorithm remains an open problem. Nevertheless, the widely documented success of the Longstaff and Schwartz (2001) methodology and its variants (Andersen and Broadie 2004, Carmona and Ludkovski 2005, Gobet et al. 2005), as well as stable numerical behavior in the examples below should be compelling empirical evidence regarding the performance of the approach. Moreover, the suggested scheme is not the only possibility, as many other methods are available to approximate the conditional expectation of (4.1). In particular, let us mention the Markov Chain approximation method (Kushner and Dupuis 2001), the optimal quantization method (Bally et al. 2005) and the kernel regression method (Györfi et al. 2002). The first two of these methods discretize the dynamics of $\{Y_t\}$ and replace them with a discrete-state Markov chain $\{\tilde{Y}_t\}$. Once this is done, conditional expectations can be computed via the standard lattice methods. The last approach evaluates (4.1) using the fully non-parametric kernel regressor rather than a pre-selected set of basis functions. Each of the above methods has its own strengths and shortcomings and full numerical comparison is beyond the scope of this paper. We chose the method of Section 4.2 for its intuitive probabilistic structure, ease of implementation and current popularity; the jury is still out whether this is the most numerically efficient and stable approach.

From a practical point of view, the main computational parameter is the number of paths N^p. Heuristically, N^p must grow exponentially in number of basis functions N^B. The dependence on Δt is unknown, but as documented above, changes in Δt have little impact on p^Δ, so Δt can be taken as fixed. The overall algorithm complexity is $\mathcal{O}(N^p \cdot (N^B)^3 \cdot \Delta t)$. Table 1 shows the standard deviation of the initial value $p_{0,T}$ as a function of N^p paths used in the simulation. We find that with $N^B = 6, \Delta t = T/364$ and $N^p = 32000$ the standard deviation in $p_{0,T}(y^0, i)$ is less than 1%, which is practically acceptable. Note that for small N^p the algorithm seems to exhibit a consistent upward bias. With the above parameters and implementation in Matlab, the running time was about five minutes on a stock office desktop.

5. Numerical Examples

To illustrate our results, we return to our oil-platform case study and consider a representative example involving a geometric Brownian motion model with three operational regimes, $i \in \{0,1,2\}$. Let

$$\left\{ \begin{array}{rcll} \mathrm{d}S_t &=& 0.05 S_t\,\mathrm{d}t + 0.4 S_t\,\mathrm{d}W^1_t, & S_0 = 50, \\ \mathrm{d}Y_t &=& 0.05 Y_t\,\mathrm{d}t + 0.4 Y_t \cdot (0.9\,\mathrm{d}W^1_t + \sqrt{1-0.9^2}\,\mathrm{d}W^2_t), & Y_0 = 50. \end{array} \right. \tag{5.1}$$

The three regimes correspond to (0) keeping the production shut down, (1) running at normal capacity of 5 million barrels/year with break-even price of \$50/bbl, and (2) running at high capacity of 10 Mbbl/yr

Number of Paths N^p	Mean $p_{0,T}(Y_0, 0)$	St. Dev. $p_{0,T}(Y_0, 0)$
4000	9.25	3.17%
8000	9.05	2.29%
16000	8.92	1.44%
24000	8.91	1.34%
32000	8.89	0.86%
40000	8.89	0.78%

TABLE 1. Mean and Standard Deviation of $p_{0,T}(Y_0, 0)$ for Example 1 below. We use six basis functions $N^B = 6$, $\Delta t = T/364$ and the algorithm of Section 4.2.

with break-even price of \$56/bbl. We assume that there are no other costs, so that the respective cashflow functions (in millions of dollars) are

$$\psi_0(y) = 0, \psi_1(y) = 5(y - 50) \wedge C_\psi, \psi_2(y) = 10(y - 56) \wedge C_\psi,$$

where the bound C_ψ is taken to be sufficiently large, e.g. $C_\psi = 2000$. The last regime is therefore preferred when prices are above \$62/bbl, $\psi_2(y) > \psi_1(y) \Leftrightarrow y > 62$. The numerical parameter values are meant to roughly correspond to oil markets typical in 2007. Finally, we assume that the switching costs are given by $C_{i,j} = 0.25|i-j|$ so that it takes \$250,000 to make a sequential regime change. The interest rate is $r = 0.05$ (so that $M_{t,T} \equiv 0$ in (3.2)) and the manager has a planning horizon of six months, $T = 0.5$. For simplicity, there are no other constraints and no price impact.

With these parameter values, the classical formula (1.1) gives $\bar{V}(Y_0) = \$12.37$ million. In contrast, taking $\gamma = 0.1$ we find

$$p_{0,T}(Y_0, 0) = 8.89, \quad p_{0,T}(Y_0, 1) = 8.86, \quad p_{0,T}(Y_0, 2) = 8.61.$$

These values were obtained by running the Monte Carlo scheme of Section 4.2 using 32000 paths, 364 time-steps and six basis functions. The resulting value for $p_{0,T}$ had the above means and a standard deviation of 0.86%.

Thus, the classical approach overestimates the true value by nearly 40%. Again, the fundamental reasons for this overestimate are the switching costs that limit flexibility and the risk-aversion of the manager. Figure 1 shows the effect of these two factors on the indifference value of production. In particular, we find that in the absence of risk-aversion ($\gamma = 0$ or equivalently $\rho = 1$), the value would be $p_{0,T}(Y_0, 0; \gamma = 0) = 11.60$, while in the absence of switching costs $C_{i,j} = 0 \,\forall i, j$, the value would be $p_{0,T}(Y_0, 0; C = 0) = 9.72$. The classical value $\bar{V}(Y_0)$ corresponds to the extreme top left corner where there are no constraints and a complete market. The top boundary $\gamma(1 - \rho^2) = 0$ of Figure 1 corresponds to a standard (risk-neutral) optimal switching problem, see Section 3.4, and was computed using the algorithm in Carmona and Ludkovski (2005). Further analysis of the effect of C and γ on the problem structure is in Sections 5.2-5.3 below.

5.1. **Optimal Policy.** Similar to American exercise boundary for early exercise options, the optimal operational policy of (2.9) can be summarized by a switching boundary plot. The switching boundary $\Gamma_{i,j}(m\Delta t)$ delineates the region where $\{p_{m\Delta t,T}(\cdot, i) = p_{m\Delta t,T}(\cdot, j) - C_{i,j}\}$ or in terms of the notation of Section 4.1,

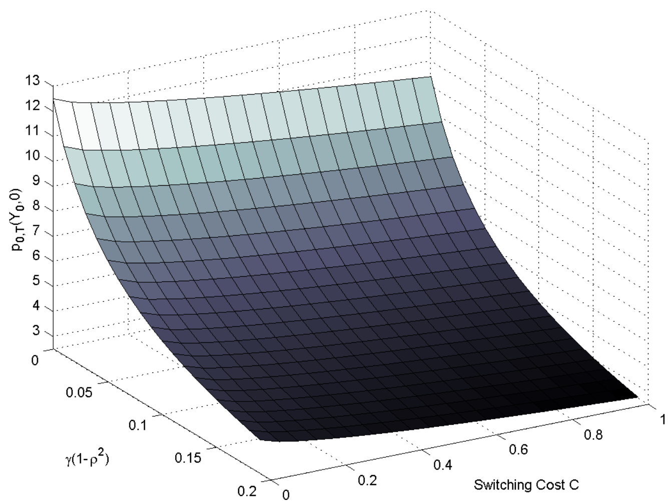

FIGURE 1. Dependence of indifference value $p_{0,T}(Y_0, 0)$ of Example 1 on the switching cost C (we take $C_{i,j} = C|i - j|$) and the market incompleteness measure $\gamma(1 - \rho^2)$.

the empirical region where $\xi^{n,i}_{m\Delta t} = j$. Figure 2 shows these switching boundaries for the preceding example. For instance, starting in the 'off' regime 0, the optimal strategy is to switch to regime 1 as soon as (and only then) the commodity price $\{Y_t\}$ reaches the lower dotted boundary $\Gamma_{0,1}(t)$ in Figure 2. In other words, $\tau_1 = \inf\{t\colon Y_t \geq \Gamma_{0,1}(t)\}$ (since the algorithm is in discrete-time, one actually only uses $t = m\Delta t$). Incidentally, the fact that $p_{0,T}(Y_0, 2) = p_{0,T}(Y_0, 1) - C_{2,1}$ indicates that starting out in the full capacity regime, the manager should immediately switch to normal capacity. Due to the natural ordering of the regimes in this specific example, it can be seen that we never switch from shutdown to full production and vice versa directly, but always go through the middle regime 1.

In the classical setting of (1.1), the switching boundaries would be straight lines at the break-even levels of $\$50$ and $\$56$ per barrel. However in our model, when the price rises just above $\$50/bbl$ and the oil-platform is shut, the manager is reluctant to start production. This is because it would entail an immediate cost $C_{0,1}$ and there is uncertainty regarding future Y-prices. Instead she will wait until prices reach about $\Gamma_{0,1}(t) \simeq \$53/bbl$ and will start production only then. Similarly, being in regime 1, the optimal strategy is to switch to regime 0 only once $\{Y_t\}$ hits the lower solid boundary $\Gamma_{1,0}(t)$ (about $\$47.5$) from above. Again, this boundary is below the break-even level of $\$50/bbl$, indicating that the manager will be willing to suffer some small losses in the hope of eventual recovery rather than immediately incur the large switching costs. Thus, the operational constraints and the managerial risk-aversion cause the appearance of the *hysteresis* band (Dixit 1989) between $\Gamma_{0,1}(t)$ and $\Gamma_{1,0}(t)$ (and similarly between $\Gamma_{1,2}(t)$ and $\Gamma_{2,1}(t)$). The hysteresis bands showcase the path-dependency of the problem and the conservative attitudes of the manager. Because the manager's behavior is time-inhomogeneous, so are the switching boundaries. In particular, close to terminal date T the hysteresis bands widen dramatically since the immediate cost of making a production mode switch dominates any possible gain to be made before T (recall that the residual value was assumed to be zero).

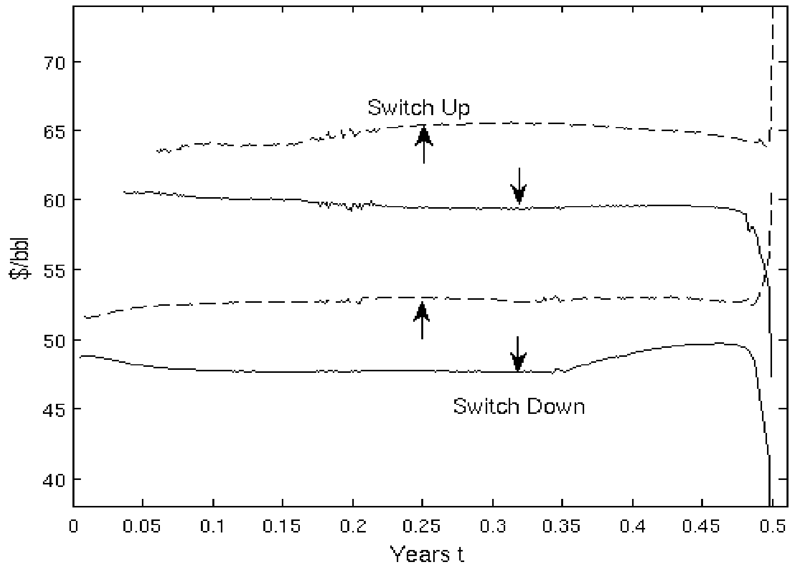

FIGURE 2. Switching boundaries for Example 1 in Section 5.1. The figure shows the four relevant boundaries $\Gamma_{1,0}(t), \Gamma_{0,1}(t), \Gamma_{2,1}(t), \Gamma_{1,2}(t)$ with $\gamma = 0.1$ as a function of time t. The dashed boundaries represent levels for switching to a higher production regime and the solid boundaries represent levels for switching to a lower production regime. Because the boundaries were generated using paths that all begin with $Y_0 = 50$, for small t none of the paths switched to regime 2 and there is no threshold to display for $\Gamma_{2,1}(t)$ and $\Gamma_{1,2}(t)$.

Once the switching boundaries are known, the optimal production strategy is completely determined. Namely, given a path of $\{Y_t\}$ one sequentially checks if Y_t crosses the appropriate switching boundaries $\Gamma_{i,j}(t)$ and makes the necessary operational regime changes. Figure 3 illustrates this procedure using the boundaries above. The top panel shows two simulated paths of $\{Y_t\}$ in relation to switching boundaries. The lower panel keeps track of cumulative realized gains $B_{0,t}(\xi^*)$ from the oil platforms, assuming initial regime $\xi^*_0 = 0$. Observe that $B_{0,t}(\xi^*)$ might decrease as the manager may be losing money, e.g. running the asset when $\Gamma_{1,0}(t) \le Y_t \le 50$. The discrete switching costs C are indicated by instantaneous drops in $B_{0,t}(\xi^*)$, see for instance the solid curve on the lower panel around $t = 0.25$. One may also compute statistics of the resulting total operational gains $B_{0,T}(\xi^*)$ under the physical measure $\mathbb{P}$. We find that $\mathbb{E}[B_{0,T}(\xi^*)] = 11.58$, and $StDev_{\mathbb{P}}(B_{0,T}(\xi^*)) = 19.57$ (also recall that $p_{0,T}(Y_0, 0) = -\frac{1}{\alpha_0} \ln \mathbb{E}[\exp(-\alpha_0 B_{0,T}(\xi^*))] = 8.89$). At the right tail we find that $\mathbb{P}(B_{0,T}(\xi^*) = 0) \simeq 0.31$ and $\mathbb{P}(B_{0,T}(\xi^*) < 0) \simeq 0.047$, so that 4.7% of the time the production will result in an overall loss and 31% of the time the oil platforms will be kept shut down throughout the six months. On the other hand, we find $\mathbb{P}(B_{0,T}(\xi^*) > 50) \simeq 0.059$, showing that highly profitable outcomes are also not uncommon.

It remains to describe the optimal hedging strategy π^*. Recall that $\pi^*_t(Y_t, i)$ denotes the dollar amount that the manager will invest in the S-asset given the current local price and the current production regime, and was computed in terms of $p_{t,T}(Y_t, i)$ in (3.13). Figure 4 shows the optimal $\pi^*_0(Y_0, 0)$ at $t = 0$ for different values of Y_0. This was obtained by recomputing $p_{0,T}(Y_0, i)$ for different initial values of Y_0 and then using a finite-difference approximation of (3.13). As expected, $\pi^*_0(Y_0, 0) < 0$ since the manager will attempt to

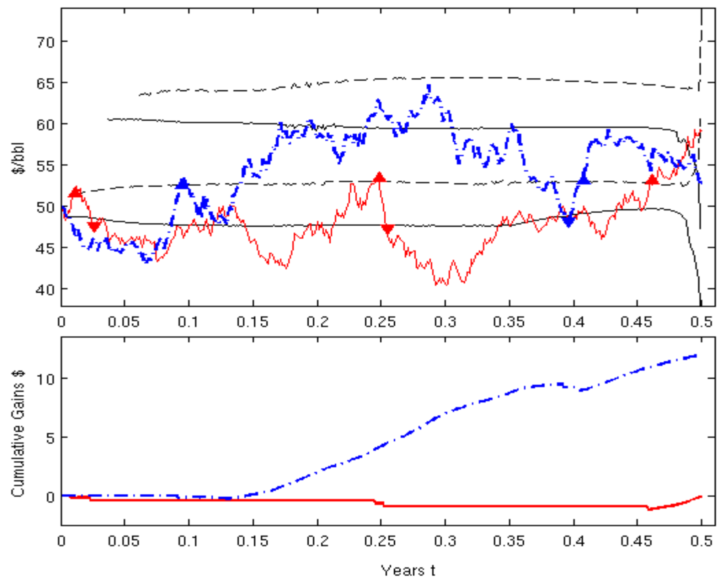

FIGURE 3. Optimal Operational Strategy for Example 1. The top panel shows evolution of two sample paths of $\{Y_t\}$, as well as the switching boundaries $\Gamma_{i,j}(t)$ of Figure 2. The bottom panel shows the corresponding cumulative realized operational gains $B_{0,t}(\xi^*)(\omega)$ as a function of t. Switching up/down times are indicated with upper and lower triangles respectively.

short the traded asset to hedge her production Call options on $\{Y_t\}$. This can also be seen from (3.13): in our example higher prices increase expected revenues so that $\partial_Y p_{0,T}(\cdot, 0) > 0$. Intuitively, the production flexibility implies that the manager has a joint option on the (continuously-paying) $Call_1 = 5(Y_t - 50)_+$ and $Call_2 = 5(2Y_t - 112)_+$. Thus, when Y_t is small, we expect $\pi_t^*(Y_t, 0) \simeq 0$, when Y_t is large we expect $\pi_t^*(Y_t, 0) \simeq \pi_t^{Call_2}$ and in between we expect $\pi_t^*(Y_t, 0) \simeq \pi_t^{Call_1}$. As Y_t increases, $\pi_t^*(Y_t, 0)$ increases in absolute value, because the Call options are deeper in-the-money. Figure 4 confirms this intuition; we see that $\pi_0^*(Y_0, 0)$ indeed interpolates between $0, \pi_0^{Call_1}$ and $\pi_0^{Call_2}$.

5.2. **Effects of Risk-Aversion and Correlation.** In our model the risk-preferences of the manager are conveniently summarized in a single parameter γ. Because the risk-aversion parameter γ and the correlation between the traded and local contracts ρ always appear together as the product $\gamma(1-\rho^2)$, the effect of changing ρ is equivalent to changing γ. We call the above product the measure of *market incompleteness*; recall that as $\gamma \to 0$ (or $\rho \to 1$), the manager becomes risk-neutral (the asset cashflow can be perfectly replicated) and we pass to the limiting case of optimal switching under the risk-neutral measure $\mathbb{Q}$. We already know from Corollary 3 that $p_{t,T}(y, i)$ is decreasing in γ (and increasing in ρ). Increase in market incompleteness can in fact be decomposed into two effects that may be termed “pessimism” and “precaution”:

- Precaution: Higher $\gamma(1-\rho^2)$ makes the manager eschew more risky regimes, further widening the hysteresis bands.
- Pessimism: Higher $\gamma(1-\rho^2)$ means that the manager is more conservative and places a lot of value on scenarios with low revenue. She will gain little utility from highly profitable outcomes.

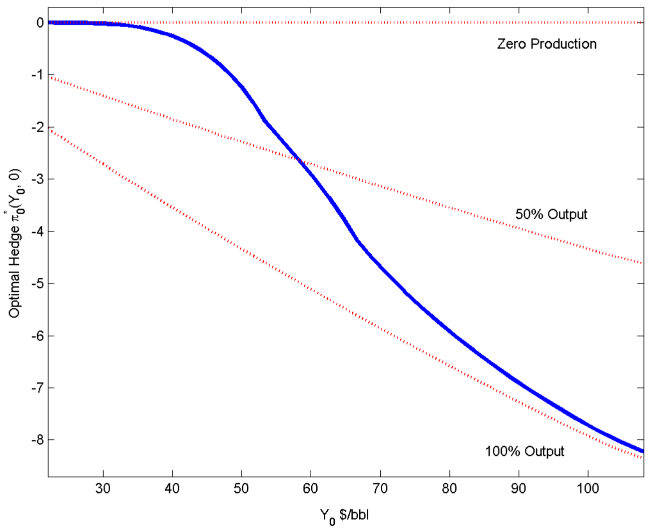

FIGURE 4. Delta hedging of operational flexibility for Example 1 at time $t = 0$. The thick blue curve shows the optimal hedge $\pi_0^*(Y_0, 0)$ at $t = 0$ for different values of Y_0 and initial regime 0. The sloping dotted lines indicate classical Delta hedging amounts $\pi_0^{Call_1}, \pi_0^{Call_2}$ computed using the equivalent of (3.13) in the absence of operational flexibility.

Figure 1 already showed the impact of changing $\gamma(1-\rho^2)$ on initial $p_{0,T}(y,i)$. The rightmost panel of Figure 5 shows how $\gamma(1-\rho^2)$ impacts the switching boundaries $\Gamma_{i,j}(t)$. Interestingly we observe that as $\gamma(1-\rho^2)$ increases, all switching boundaries *increase*, with a very slight widening of the hysteresis bands. This is because the volatility of revenues is highest in regime 2 and lowest in regime 0. Thus, precaution encourages the manager to spend more time in the least-volatile regime 0, increasing both the up-switching boundary $\Gamma_{0,1}(t)$ and the down-switching boundary $\Gamma_{1,0}(t)$ (similar effect for $\Gamma_{1,2}(t)$ and $\Gamma_{2,1}(t)$). This slight change should be contrasted with the strong pessimism effect manifested in Figure 1.

Remark 6. The simulation algorithm is sensitive to the product $\gamma(1-\rho^2)$ which enters the power in the nonlinear expectation (3.15). Thus the variance of the algorithm increases as $\gamma(1-\rho^2)$ decreases.

5.3. **Effect of Switching Costs.** As stated in Corollary 3, the indifference value $p_{0,T}(Y_0, 0)$ is decreasing in switching costs $C_{i,j}$. As $C_{i,j} \to 0$, the path-dependency of the problem disappears and the current regime has no influence. Thus, at every switching opportunity, the manager will simply choose the regime with the highest payoff, so that her flexibility becomes a series of chooser Call options. Accordingly, the hysteresis bands shrink away. The middle panel of Figure 5 demonstrates this feature by plotting $\Gamma_{i,j}(0.2)$ against switching cost scale C, where we have taken $C_{i,j} = C|i-j|$. At the other extreme, high $C_{i,j}$ makes changing regimes very expensive and takes away much of the managerial flexibility ($C = 1$ takes away nearly 20% of value).

5.4. **Effect of Volatility.** As the volatility of local price b increases, two conflicting events take place: (a) ceteris paribus switching becomes more frequent, increasing switching costs; (b) there are more upside opportunities which bring in extra profit since the manager has Call options on income. In this particular

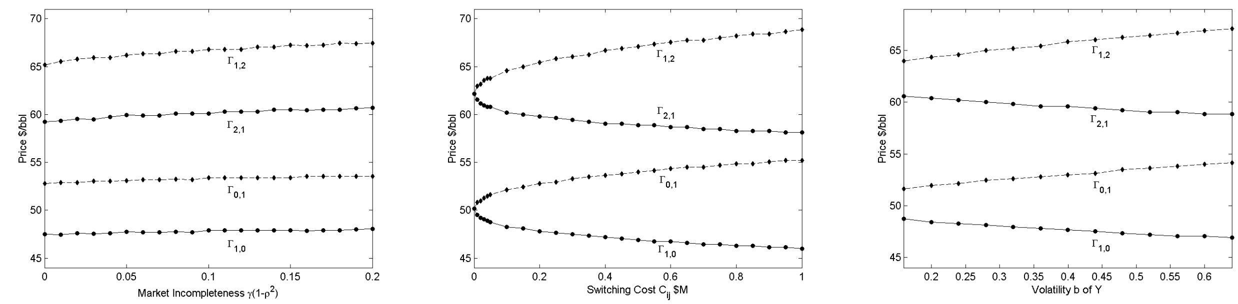

FIGURE 5. Dependence of switching boundaries $\Gamma_{i,j}(t)$ of Example 1 for a fixed $t = 0.2$ on model parameters. From left-to-right: the effect of the measure of market incompleteness $\gamma(1-\rho^2)$; the effect of switching costs C in $C_{i,j} = C|i-j|$; the effect of production-price volatility b.

example, we find that the second effect dominates and the first one is mitigated by the widening of the hysteresis bands, see the rightmost panel of Figure 5. This happens because of precaution: there is more uncertainty about future evolution of $\{Y_t\}$, so each production switch is more risky and taken more reluctantly. Overall, the indifference value $p_{0,T}(Y_0, 0)$ is slightly convex in b, a behavior similar to standard Call options. For instance, compared to $b = 0.4$, $b = 0.32$ reduces value by 24%; with $b = 0.48$, $p_{0,T}(Y_0, 0)$ is increased by 25%.

5.5. **Further features: Example 2.** To further illustrate the capabilities and structure of our model, we consider a more complicated second example. It features the manager of a gold mining firm who has four possibilities $\xi_t \in \{0, 1, 2, 3\}$ regarding running a particular mine location:

- Mothball the site which carries zero costs, $\psi_0(Y_t) \equiv 0$;
- Temporary shutdown which carries fixed costs of $K = -\$40M/yr$, $\psi_1(Y_t) = -40$;
- Normal operation with extraction costs of \$ 530/ounce and production rate of 1 million ounces a year, plus the aforementioned fixed costs, $\psi_2(Y_t) = (Y_t - 530) - 40$;
- Maximum production of 1.5 million ounces a year but with a a slightly higher fixed cost of \$55 million a year to reflect hiring of extra labor: $\psi_3(Y_t) = 1.5Y_t - 850$;
- The cost matrix is given by

$$C = (C_{i,j}) = \begin{pmatrix} 0 & 25 & 25 & 50 \\ 25 & 0 & 5 & 25 \\ 25 & 0 & 0 & 25 \\ 50 & 25 & 25 & 0 \end{pmatrix},$$

  so that mothballing and maximum production are very expensive to initiate and end;
- The firm makes twice-weekly decisions about its operations so the operational flexibility problem has a given $\Delta t = 1/104$; the planning horizon is $T = 2$ years.
- All decisions take a full two weeks to implement, during which no further operational switches are possible.
- The company hedges its production using the continuously-traded liquid London bullion market $\{S_t\}$ where $\mathrm{d}S_t = S_t(0.05\,\mathrm{d}t + 0.2\,\mathrm{d}W_t^1)$.

- Production is tied to the local wholesale price which is modeled by

$$\mathrm{d}Y_t^{\xi} = \kappa(\ln\theta_Y(\xi_t) - \ln Y_t^{\xi}) \cdot Y_t^{\xi}\,\mathrm{d}t + bY_t^{\xi} \cdot (0.99\,\mathrm{d}W_t^1 + \sqrt{1-0.99^2}\,\mathrm{d}W_t^2). \tag{5.2}$$

- Above parameters are $\theta_Y(0) = \theta_Y(1) = \theta_Y(2) = 600$, $\theta_Y(3) = 580$, $\kappa = 1$, $b = 0.2$, $r = 0.06$.
- The risk-aversion coefficient is $\gamma = 0.05$.

The dynamics (5.2) is the log-normal mean-reverting model often used for commodities (Schwartz 1997). For a fixed production regime, $\{Y_t^i\}$ will tend to be around $\theta_Y(i)$, more precisely $\{\ln Y_t^i\}$ follows an Ornstein-Uhlenbeck process with mean-reversion level $\ln\theta_Y(i)$. Observe that this level is taken to depend on ξ_t, modeling *price impact*. Namely we suppose that when extraction proceeds at maximum rate, local supply of gold increases and drives prices down from a mean of $\$600/oz$ to $\$580/oz$.

The operational inertia feature described implies that $\delta_i = 4\Delta t$ in the notation of Remark 2. This is incorporated into the numerical algorithm by adjusting (3.15) to

$$p^{\Delta}_{t,T}(y,i) = \left\{ \max_{j\neq i} \frac{-1}{\alpha_t} \ln \bar{\mathbb{E}}_{t,y,i}\Big[\exp\big(-\alpha_{t+\delta_j}[p^{\Delta}_{t+\delta_j,T}(Y_{t+\delta_j}, j) + B_{t,t+\delta_j}(j) - \mathrm{e}^{r\delta_j}C_{i,j}]\big)\Big]\right\}$$
$$\vee \left\{ \frac{-1}{\alpha_t} \ln \bar{\mathbb{E}}_{t,y,i}\Big[\exp\big(-\alpha_{t+\Delta t}[p^{\Delta}_{t+\Delta t,T}(Y_{t+\Delta t}, i) + B_{t,t+\Delta t}(i)]\big)\Big]\right\}.$$

Observe that in the model considered, the measure of market incompleteness is rather small $\gamma(1-\rho^2) = 9.95\cdot 10^{-4}$ and the switching costs are relatively high. Consequently, the precaution effect is expected to be much more influential compared to Example 1. Also, the manager must balance between temporary shutdowns (regime 1) that maintain future flexibility, and mothballing that eliminates the fixed costs K. We find that for $Y_0 = \$600/oz$,

$$p^{\Delta}_{0,T}(Y_0, \cdot) = [45.06, \quad 55.33, \quad 61.58, \quad 60.45].$$

For comparison, $\bar{\mathbb{E}}[\int_0^2 (Y_t - 570)\,\mathrm{d}t\,|Y_0 = 600] = 47.13$, so even though the mine is generally profitable and one expects regime 2 to be most common, the flexibility of the manager increases value by over \$14 million, or almost 30% compared to the case of just running it in standard production throughout the two years.

As illustrated in Figure 6 the structure of the switching boundaries in this example is rather non-trivial:

- From mothballed regime 0, the decision maker will wait until prices rise all the way to $\Gamma_{0,2}(t) \simeq \$640/oz$ and then switch directly into regular production;
- From the 'off' regime, the manager will either mothball the mine if prices drop to about $\Gamma_{1,0}(t) \simeq \$480/oz$ (recall that regular production cost is $\$530/oz$), or go to regular production if prices rise to about $\Gamma_{1,2}(t) \simeq \$550/oz$;
- From regular production, the firm will either go into no production if price drops to $\Gamma_{2,1}(t) \simeq \$495/oz$, or into maximum production if prices skyrocket to $\Gamma_{2,3}(t) \simeq \$725/oz$. Observe that even though the dis-economy of scale is very small, the price impact and high switching costs make the manager reluctant to go into maximum production.
- Finally, from maximum production the firm will keep producing until prices drop to $\max(\Gamma_{3,2}(t), \Gamma_{3,1}(t)) \simeq \$495/oz$. Early on it will then shut down production, however after a few months the best action becomes switching to regular production instead. This bifurcation shows that the time-dependence can lead to quite complex effects.

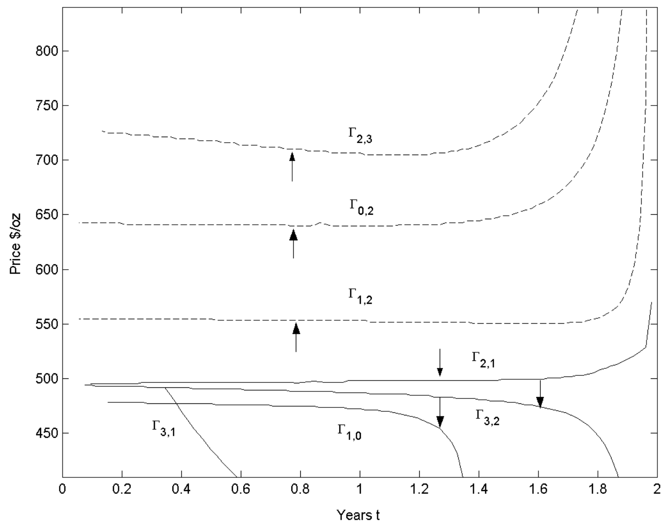

FIGURE 6. Switching boundaries for Example 2 of Section (5.5). We show the seven relevant boundaries $\Gamma_{0,2}(t), \Gamma_{1,2}(t), \Gamma_{2,3}(t), \Gamma_{3,2}(t), \Gamma_{3,1}(t), \Gamma_{2,1}(t), \Gamma_{1,0}(t)$ as a function of time t. The dashed boundaries with upward arrows represent levels for switching to a higher production regime and the solid boundaries and downward arrows represent levels for switching to a lower production regime.

5.6. **Effect of Operational Constraints.** We now investigate the effect of changing the various constraints faced by the gold producer. The results are summarized in Table 2. As the first step, we study the effect of operational delay δ. One could imagine that the firm can invest in a more streamlined execution structure that will reduce δ and wants to assess the resulting benefits to determine the viability of such an upgrade. The first few rows of Table 2 show the effect of doing so for various values of δ, compared to the base case of $\delta = 4\Delta t$. While the absolute changes are not very large, compared to the whole "flexibility" benefit of $14 million they are significant. Thus, going to $\delta = \Delta t$ (so that a decision leads to effective production change within half a week) will increase flexibility benefit by another 14%. For comparison, if implementation delay were to become a full month, almost $2 millon would be lost.

We next consider the extra flexibility afforded by the mothballing regime 0. First, imagine that mine mothballing must be permanent, which is equivalent to setting $C_{0,j} = +\infty$ for all j. We find that the associated value drops by $\$0.2$ million, which is very small. On the other hand, if we imagine that mothballing is not possible at all (equivalent to $C_{j,0} = +\infty$ for all j), then value drops by another $\$0.6$ million or a total loss of 6% of flexibility benefit. Thus, we see that regime 0 does make a significant contribution to overall flexibility of the manager and that it is indeed used as a long-term mothballing state (since one very rarely switches out of it).

Finally, we can consider by how much $p_{0,T}(y,i)$ is reduced by the price impact. Changing to $\theta_Y(3) = 600$ we find that $p_{0,T}(600,1)$ increases by nearly $7 million, showing that market power can have very strong effect on overall profitability. The effect of ξ on $\{Y_t\}$ also explains the extreme reluctance of the manager to switch into high production ($\Gamma_{2,3}(t)$ of over $\$720/oz$ in Figure 6).

Model Change	$p_{0,T}(Y_0, 1)$	$p_{0,T}(Y_0, 2)$
Base Case $\delta = 4\Delta t$	55.33	61.58
$\delta = 8\Delta t$	53.29	60.73
$\delta = 2\Delta t$	56.62	62.26
$\delta = \Delta t$	57.31	62.64
Permanent Mothballing	55.08	61.32
No Mothballing Possible	54.49	60.68
No Price Impact $\theta_Y(3) = 600$	62.23	68.78

TABLE 2. Effect of various operational settings on the indifference price of operational flexibility in Example 2 of Section 5.5. Values were computed using the simulation algorithm of Section 4.2 with $N^p = 40,000$ paths and $N^B = 6$.

6. CONCLUSION

In this paper we have studied the problem of optimal firm management using a joint operational/financial strategy. This approach allows overall risk-management of the firm with both financial and operational policies given equal footing. Our model links the methods of option pricing in incomplete markets (Carmona, ed. 2006) with the real options literature (Dixit and Pindyck 1994). This is reflected in the solution structure which consists of standard portfolio optimization problems between production decisions and switching boundaries that determine production regimes. Moreover, the representation as a series of coupled American options is intuitive for the manager and gives simple and familiar policy guidelines. Besides being comprehensive, as illustrated in the last section, our model also permits a highly granular approach for analyzing the effects of various operational constrains, ranging from switching costs to impact of market power. Looking forward, our model can also be used for strategic production planning (expansion, mergers, upgrades, etc.) via comparison of indifference values $p_{t,T}(y, i; \mathcal{U})$ for assets with various acceptable policy sets $\mathcal{U}$.

In terms of model parameters, we found that the risk-aversion of the manager induces a strong pessimism effect (on the order of 10%-50%) on the value $p_{0,T}(y, i)$ of the asset, and a weak precaution effect (on the order of 1%-5% in terms of $\Gamma_{i,j}$) on the optimal production policy ξ^*. On the contrary, the fixed switching costs have only a $2\% - 5\%$ effect on $p_{0,T}(y, i)$ but strongly influence the width of the hysteresis bands. Since the optimal operational policy ξ^* is determined by $\{\Gamma_{i,j}(t)\}$, a possible shortcut for approximating the full model (2.9) is to solve a standard linear optimal switching problem under the minimal martingale measure $\mathbb{Q}$. This will (nearly fully) reflect the effect of switching constraints; the pessimism effect of $\gamma(1 - \rho^2)$ can then be added-on via a volatility penalty on $B_{0,T}(\xi^*)$, as explained by Henderson (2002):

$$-\frac{1}{\alpha_t} \ln \bar{\mathbb{E}}_{t,y,i} \left[\exp(-\alpha_T B_{t,T}(\xi^*)) \right] \simeq e^{-r(T-t)} \left\{ \bar{\mathbb{E}}_{t,y,i}[B_{t,T}(\xi^*)] - \frac{\gamma(1-\rho^2)}{2} Var_{t,y,i}(B_{0,T}(\xi^*)) + \mathcal{O}\left(\gamma^2(1-\rho^2)^2\right) \right\}.$$

6.1. **Liquidity Risk.**

A different interesting application of our framework can be found in the area of liquidity risk. Consider a portfolio optimization problem using one liquid asset with negligible transaction costs, and one illiquid asset with significant transaction costs. Such a situation arises whenever a thinly traded underlying $\{Y_t\}$ (e.g. an exotic commodity forward) is tracked using a closely related market contract $\{S_t\}$. The action set of ξ would now correspond to various possible positions taken in Y and the functions $\psi_i(t,y)$ will reflect the resulting P&L. The fixed switching costs $C_{i,j}$ model the liquidity constraints of trading in $\{Y_t\}$, and one can also incorporate again price impact of the trading strategy ξ on the illiquid price $\{Y_t^\xi\}$. A related model has been recently studied by Ly Vath et al. (2007) using analytic pde methods.

6.2. **General Utility Processes.** Recall that in the case of exponential utility we obtained *iterativity* of the indifference value in (3.12) across time-periods. This natural feature is very attractive and matches one of the basic properties of the classical method (1.1). We stress that it is not intrinsic to the indifference valuation method; rather it follows from our choice of exponential utility $U(x)$. In fact, as shown by Cheridito and Kupper (2006), within the expected utility framework, exponential utility is the only one that satisfies this iterative structure of (3.12).

However, it is possible to extend the model beyond terminal expected utility framework. Note that the latter is actually rather limited in practice, as it assumes the manager only cares about *final* wealth level at T. A realistic manager is likely to have preferences not over wealth levels at date T, but over entire wealth-paths over $[0,T]$. For instance, the firm might prefer to avoid having the wealth dip too low at intermediate time points to prevent a credit crunch, or it might wish to do earnings management by minimizing volatility of cashflows. Such a setup can be accommodated in the framework of monetary utility processes introduced by Cheridito et al. (2006). The idea is to replace the single utility function $U(X_T)$ with a sequence of utility functionals of the form $U_{t,T}(X_\cdot)$, which assign to a wealth process $(X_t)_{t\le T}$ its risk-adjusted date-t value based on the entire evolution from t to T. Thus, $U_{t,T}$ is a map from say $L^\infty([t,T]\times\mathbb{R})$ to $L^\infty(\mathcal{F}_t)$. In this paper, we worked with the entropic monetary utility process where

$$U_{t,T}(X_\cdot) = -\frac{1}{\gamma}\ln \mathbb{E}^{\mathbb{Q}}\big[\exp(-\gamma X_T)\,|\,\widetilde{\mathcal{F}}_t\big]. \tag{6.1}$$

Another popular possibility is the worst-stopping functional $U_{t,T}(X_\cdot) = \operatorname{ess\,inf}_{\tau\in\mathcal{S}(t)}\mathbb{E}\big[X_\tau\,|\,\mathcal{F}_t\big]$. Once a particular family $\{U_{t,T}\}_{t\le T}$ is picked, the optimization problem (2.9) becomes

$$V(t,x,y,S,i) = \operatorname*{ess\,sup}_{\pi\in\mathcal{A}(t,T),\xi\in\mathcal{U}(t,T)} U_{t,T}\big(X_\cdot^{t,x,\pi,\xi}\big), \qquad Y_t=y, S_t=s, \xi_t=i$$

and the resulting indifference valuation method is

$$p^U_{t,T}(y,i) = \operatorname*{ess\,sup}_{\pi\in\mathcal{A}(t,T),\xi\in\mathcal{U}(t,T)} U_{t,T}\big(X_\cdot^{t,x,\pi,\xi}\big) - \operatorname*{ess\,sup}_{\pi\in\mathcal{A}(t,T)} U_{t,T}\big(X_\cdot^{t,x,\pi}\big), \qquad \text{with } Y_t=y, \xi_t=i. \tag{6.2}$$

For this to make economic sense, it is necessary that the family $\{U_{t,T}\}$ be *time-consistent* (Cheridito et al. 2006). Time-consistency is also sufficient for the problem to be invariant with respect to initial wealth x and for $p_{t,T}$ to satisfy a recursion similar (3.12). Computationally, both properties are crucial for tractability and also allow one to use the Longstaff and Schwartz (2001) approach of Section 4.2 over the coupled optimal stopping problems. The theory of monetary utility processes in continuous-time is still incomplete (see some recent progress in Klöppel and Schweizer (2007)), but offers exciting possibilities for more realistic modeling of our problem. We leave further work in this direction to future research.

Another alternative is to consider a horizon-independent construction which removes T altogether. Superficially this would rule out an expected utility framework. However, thanks to time-consistency of exponential utility, this is in fact possible using the ideas of Henderson and Hobson (2007). Unfortunately, it is not clear how to choose a proper risk criterion over operational production, in particular given seasonality in commodity markets that necessitate keeping track of calendar time and exclude time-stationary solutions.

APPENDIX: PROOFS

Proof of Lemma 3. We prove the lemma by induction. First for $k = 1$, the control set is simply $\mathcal{U}^1(\sigma, T) = \{\xi = (\sigma, \xi_0, \tau_1, \xi_1)\}$ where the only choices are for $\tau_1 \in \mathcal{S}(\sigma)$ and $\xi_1 \in \mathcal{D}(\tau_1, \xi_0)$. Hence, using (2.6) and $\alpha_{\tau_1} = \alpha_T \mathrm{e}^{r(T-\tau_1)}$,

$$\begin{aligned}
\phi^1(\sigma, Y_\sigma, i) &= \operatorname*{ess\,sup}_{\xi \in \mathcal{U}^1(\sigma,T), \xi_\sigma = i} \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\sigma,T}(\xi)) \Big| \mathcal{F}_\sigma \right] \\
&= \operatorname*{ess\,sup}_{\tau_1 \in \mathcal{S}(\sigma), \xi_1 \in \mathcal{D}(\tau_1, i)} \bar{\mathbb{E}}\left[-\exp\Big(-\alpha_T [\mathrm{e}^{r(T-\tau_1)}(B_{\sigma,\tau_1}(i) - C_{i,\xi_1}) + B_{\tau_1,T}(\xi_1)]\Big) \Big| \mathcal{F}_\sigma \right] \\
&= \operatorname*{ess\,sup}_{\tau_1 \in \mathcal{S}(\sigma), \xi_1 \in \mathcal{D}(\tau_1, i)} \bar{\mathbb{E}}\left[\exp(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]) \cdot \bar{\mathbb{E}}\big[-\exp(-\alpha_T B_{\tau_1,T}(\xi_1)) \big| \mathcal{F}_{\tau_1} \big] \Big| \mathcal{F}_\sigma \right] \\
&= \operatorname*{ess\,sup}_{\tau_1 \in \mathcal{S}(\sigma), \xi_1 \in \mathcal{D}(\tau_1, i)} \bar{\mathbb{E}}\left[\exp(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]) \cdot \phi^0(\tau_1, Y_{\tau_1}, \xi_1) \Big| \mathcal{F}_\sigma \right].
\end{aligned}$$

Next, suppose (3.9) is true for k and consider (3.9) when k is replaced by $k+1$. Let $\xi \in \mathcal{U}^{k+1}(\sigma, T)$ be an arbitrary control with $\xi_\sigma = i$. Writing $\xi = (\sigma, i, \tau_1, \xi_1, \xi^{(k)})$ with $\xi^{(k)} \in \mathcal{U}^k(\tau_1, T)$ we obtain

$$\begin{aligned}
&\bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\sigma,T}(\xi)) \Big| \mathcal{F}_\sigma \right] = \bar{\mathbb{E}}\left[-\exp\Big(-\alpha_T \big[\mathrm{e}^{r(T-\tau_1)}(B_{\sigma,\tau_1}(i) - C_{i,\xi_1}) + B_{\tau_1,T}(\xi^{(k)})\big]\Big) \Big| \mathcal{F}_\sigma \right] \\
&\quad = \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \bar{\mathbb{E}}\left[-\exp\big(-\alpha_T B_{\tau_1,T}(\xi^{(k)})\big) \big| \mathcal{F}_{\tau_1}, \xi^{(k)}_{\tau_1} = \xi_1 \right] \Big| \mathcal{F}_\sigma \right] \\
&\quad \le \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1) \Big| \mathcal{F}_\sigma \right],
\end{aligned}$$

where the last line is by the induction hypothesis. Taking essential supremum (with respect to ξ on the left-hand-side and with respect to (τ_1, ξ_1) on the right-hand-side) yields

$$\begin{aligned}
\phi^{k+1}(\sigma, Y_\sigma, i) &= \operatorname*{ess\,sup}_{\xi \in \mathcal{U}^{k+1}(\sigma,T), \xi_\sigma = i} \bar{\mathbb{E}}\left[-\exp\big(-\alpha_T B_{\sigma,T}(\xi)\big) \Big| \mathcal{F}_\sigma \right] \\
&\le \operatorname*{ess\,sup}_{\tau_1 \in \mathcal{S}(\sigma), \xi_1 \in \mathcal{D}(\tau_1, i)} \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1) \Big| \mathcal{F}_\sigma \right]. \qquad \text{(A.1)}
\end{aligned}$$

Conversely, let (τ_1, ξ_1) be an arbitrary pair in $(\mathcal{S}(\sigma), \mathcal{D}(\tau_1, i))$. It is easy to check that the control set $\mathcal{U}^k(\tau_1, T)$ is stable under pairwise maximization (directed upwards) and therefore there exists a sequence $\xi^{(n)} \in \mathcal{U}^k(\tau_1, T)$, $\xi^{(n)}_{\tau_1} = \xi_1$ such that

$$\phi^k(\tau_1, Y_{\tau_1}, \xi_1) = \lim_{n\to\infty} \uparrow \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\tau_1,T}(\xi^{(n)})) \big| \mathcal{F}_{\tau_1} \right].$$

Then, for $\hat{\xi}^{(n)} \triangleq (\sigma, i, \tau_1, \xi_1, \xi^{(n)}) \in \mathcal{U}^{k+1}(\sigma, T)$ we have

$$\begin{aligned}
\phi^{k+1}(\sigma, Y_\sigma, i) &\ge \limsup_{n\to\infty} \bar{\mathbb{E}}\left[-\exp\Big(-\alpha_T B_{\sigma,T}(\hat{\xi}^{(n)})\Big) \big| \mathcal{F}_\sigma \right] \\
&= \limsup_{n\to\infty} \bar{\mathbb{E}}\left[-\exp\Big(-\alpha_T \big[\mathrm{e}^{r(T-\tau_1)}(B_{\sigma,\tau_1}(i) - C_{i,\xi_1}) + B_{\tau_1,T}(\xi^{(n)})\big]\Big) \Big| \mathcal{F}_\sigma \right] \\
&= \limsup_{n\to\infty} \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\tau_1,T}(\xi^{(n)})) \big| \mathcal{F}_{\tau_1} \right] \Big| \mathcal{F}_\sigma \right]
\end{aligned}$$

by Monotone Convergence Theorem the limit can be passed inside the expectation

$$= \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \lim_{n\to\infty} \bar{\mathbb{E}}\left[-\exp(-\alpha_T B_{\tau_1,T}(\xi^{(n)}))\Big|\, \mathcal{F}_{\tau_1}\right] \Big|\, \mathcal{F}_\sigma\right]$$
$$= \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1)\Big|\, \mathcal{F}_\sigma\right].$$

Since τ_1 and ξ_1 were arbitrary, taking essential supremum on the right hand side shows that

$$\phi^{k+1}(\sigma, Y_\sigma, i) \geq \operatorname*{ess\,sup}_{\tau_1\in\mathcal{S}(\sigma),\xi_1\in\mathcal{D}(\tau_1,i)} \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}[B_{\sigma,\tau_1}(i) - C_{i,\xi_1}]\big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1)\Big|\, \mathcal{F}_\sigma\right],$$

and combined with (A.1) concludes the proof of Lemma 3. □

Proof of Lemma 4. The proof will be established by induction. For $k = 0$, take

$$\bar{\phi}^{0,i}_t = \bar{\mathbb{E}}\left[-\exp(-\gamma(1-\rho^2)\int_0^T \mathrm{e}^{r(T-s)}\psi_i(s, Y_s)\,\mathrm{d}s)\Big|\, \mathcal{F}_t\right] \cdot \exp\left(\gamma(1-\rho^2)\int_0^t \mathrm{e}^{r(T-s)}\psi_i(s, Y_s)\,\mathrm{d}s\right).$$

Then $\bar{\phi}^{0,i}$ is a product of two smooth functionals of the Feller process Y and clearly $\phi^0(\sigma, Y_\sigma, i) = \bar{\phi}^{0,i}_\sigma$.

Now suppose that the lemma has been proved for k, and consider the case where k is replaced by $k+1$. Fix $i \in \mathbb{Z}_I$ and define

$$\check{\phi}^{k,i}_t \triangleq \max_{j\neq i}\Big\{\exp(-\alpha_t(B_{0,t}(i) - C_{i,j})) \cdot \phi^k(t, Y_t, j)\Big\}.$$

By the induction hypothesis, $\check{\phi}^{k,i}$ is a continuous, $\mathcal{F}$-progressively measurable, bounded process. General theory (El Karoui 1981) then implies that the Snell envelope $\tilde{\phi}^{k+1,i}(\sigma) = \operatorname{ess\,sup}_{\tau\in\mathcal{S}(\sigma)} \bar{\mathbb{E}}\big[\check{\phi}^{k,i}_\tau\big|\, \mathcal{F}_\sigma\big]$ of $\check{\phi}^{k,i}$ is a regular $\mathcal{F}$-supermartingale, i.e. there is a continuous $\mathcal{F}$-adapted process $\bar{\phi}^{k+1,i}$ such that $\tilde{\phi}^{k+1,i}(\sigma) = \bar{\phi}^{k+1,i}_\sigma$. Furthermore, the optimal stopping problem $\sup_{\tau\in\mathcal{S}} \bar{\mathbb{E}}\big[\check{\phi}^{k,i}_\tau\big]$ has an optimal solution explicitly given by $\tau^* = \inf\{s\colon \check{\phi}^{k,i}_s = \bar{\phi}^{k+1,i}_s\}$. We now compute

$$\begin{aligned}
\phi^{k+1}(\sigma, Y_\sigma, i) &= \operatorname*{ess\,sup}_{\tau_1\in\mathcal{S}(\sigma),\xi_1\in\mathcal{D}(\tau_1,i)} \bar{\mathbb{E}}\left[\exp(-\alpha_{\tau_1}(B_{\sigma,\tau_1}(i) - C_{i,\xi_1})) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1)\Big|\, \mathcal{F}_\sigma\right] \\
&= \operatorname*{ess\,sup}_{\substack{\tau_1\in\mathcal{S}(\sigma),\\ \xi_1\in\mathcal{D}(\tau_1,i)}} \bar{\mathbb{E}}\left[\exp\big(-\alpha_{\tau_1}(B_{0,\tau_1}(i) - \int_0^\sigma \mathrm{e}^{r(\tau_1-s)}\psi_i(s, Y_s)\,\mathrm{d}s - C_{i,\xi_1})\big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1)\Big|\, \mathcal{F}_\sigma\right] \\
&= \exp(\alpha_\sigma B_{0,\sigma}(i)) \cdot \tilde{\phi}^{k+1,i}(\sigma) = \exp(\alpha_\sigma B_{0,\sigma}(i)) \cdot \bar{\phi}^{k+1,i}_\sigma.
\end{aligned}$$

On the last line as a function of σ we have a continuous *process*, which establishes the required regularity of $\phi^{k+1}(\cdot,\cdot,i)$. The uniform bounds on $\ln(-\phi^{k+1})$ easily follow from the bounded payoff rates:

$$|\ln -\phi^k(t, y, i)| \leq \ln \bar{\mathbb{E}}[\exp(\gamma(1-\rho^2)\int_t^T \mathrm{e}^{r(T-t)}C_\psi\,\mathrm{d}s)] \leq \gamma(1-\rho^2)(T-t)\mathrm{e}^{r(T-t)}C_\psi.$$

□

Proof of Lemma 5. This follows by an easy induction argument. Indeed, the equality for $k = 0$ holds by definition of p^0 and assuming it is true for k we have from (3.9)

$$\begin{aligned}
\phi^{k+1}(t, y, i) &= \sup_{\tau_1\in\mathcal{S}(t),\xi_1\in\mathcal{D}(\tau_1,i)} \bar{\mathbb{E}}\left[\exp\Big(-\alpha_{\tau_1}[B_{t,\tau_1}(i) - C_{i,\xi_1}]\Big) \cdot \phi^k(\tau_1, Y_{\tau_1}, \xi_1)\Big| Y_t = y\right] \\
&= \sup_{\tau_1\in\mathcal{S}(t),\xi_1\in\mathcal{D}(\tau_1,i)} \bar{\mathbb{E}}\left[-\exp\Big(-\alpha_{\tau_1}\Big\{B_{t,\tau_1}(i) - C_{i,\xi_1} + p^k(\tau_1, Y_{\tau_1}, \xi_1)\Big\}\Big)\Big| Y_t = y\right] \\
&= -\exp(-\alpha_t \cdot \mathcal{G}p^k(t, y, i)) = -\exp(-\alpha_t \cdot p^{k+1}(t, y, i)).
\end{aligned}$$

by using the induction hypothesis for the second equality and (3.11) for the third one. As proved in Lemma 4, $|\ln(-\phi^k(t,y,i))|$ is bounded, which implies the same for $|p^k(t,y,i)| \le \gamma(1-\rho^2)(T-t)\mathrm{e}^{r(T-t)}C_\psi$. $\square$

Proof of Proposition 2. By Lemma 5 we have that $p^k(t,y,i)$ is increasing in k since the corresponding control sets in (3.8) are growing. On the other hand, since $\mathcal{U}^k \subset \mathcal{U}$, $p^k(t,y,i) \leqslant p_{t,T}(y,i) < \infty$, and the pointwise limit $p^\infty = \lim_{k\to\infty} p^k$ is well-defined and finite. Clearly $p^\infty(t,y,i) \leqslant p_{t,T}(y,i)$. To show that $p^\infty(t,y,i) \geqslant p_{t,T}(y,i)$, it suffices to show that for any $\epsilon > 0$, one can find a ϵ-optimal policy of $p_{t,T}(\cdot)$ which is finite, i.e. belongs to some $\mathcal{U}^K(t,T)$ for K large enough.

Let ξ^ϵ be an ϵ-optimal policy of $p_{t,T}(y,i)$. Since ξ^ϵ is admissible, $\tau_k^\epsilon \to T$ in probability, and there is a K large enough so that $\mathbb{Q}_{t,y,i}(\tau_K < T-\epsilon) < \epsilon$. For that K, take $\xi^K \in \mathcal{U}^K$ to match ξ^ϵ up to the K-th switch, with no switches after τ_K^ϵ: $\xi^K(t) = \xi^\epsilon(t)\mathbb{1}_{t<\tau_K^\epsilon} + \xi^\epsilon(\tau_K^\epsilon)\mathbb{1}_{t\ge\tau_K^\epsilon}$. Then using the fact that operational payoffs are bounded, and letting $A = \{\tau_K^\epsilon \ge T-\epsilon\}$, we have for any strategy ξ

$$|\alpha_T B_{\tau_K^\epsilon,T}(\xi)| \le \alpha_T C_\psi T \mathrm{e}^{rT}\mathbb{1}_{A^c} + \epsilon C_\psi \mathbb{1}_A =: C_1\mathbb{1}_{A^c} + C_2\epsilon\mathbb{1}_A,$$

for some constants C_1, C_2 independent of ξ. Applying the above, relation (2.6), $\mathbb{Q}_{t,y,i}(A) > 1-\epsilon$ and using the fact that $|\alpha_{\tau_K^\epsilon} B_{t,\tau_K^\epsilon}(\xi^\epsilon)| \le \gamma(1-\rho^2)(T-t)\mathrm{e}^{r(T-t)}C_\psi =: C_3$ we have

$$\begin{aligned}
p^K(t,y,i) - (p_{t,T}(y,i)-\epsilon) &\ge p^K(t,y,i;\xi^K) - p_{t,T}(y,i;\xi^\epsilon) \\
&= \frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_T(\mathrm{e}^{r(T-\tau_K^\epsilon)}B_{t,\tau_K^\epsilon}(\xi^\epsilon) + B_{\tau_K^\epsilon,T}(\xi^\epsilon)))\big]}{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_T(\mathrm{e}^{r(T-\tau_K^\epsilon)}B_{t,\tau_K^\epsilon}(\xi^\epsilon) + B_{\tau_K^\epsilon,T}(\xi^\epsilon_{\tau_K^\epsilon})))\big]} \\
&\ge \frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_{\tau_K^\epsilon}B_{t,\tau_K^\epsilon}(\xi^\epsilon))\cdot\{\exp(-C_1\mathbb{1}_{A^c} - C_2\epsilon\mathbb{1}_A)\}\big]}{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_{\tau_K^\epsilon}B_{t,\tau_K^\epsilon}(\xi^\epsilon))\cdot\{\exp(C_1\mathbb{1}_{A^c} + C_2\epsilon\mathbb{1}_A)\}\big]} \\
&\ge \frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_{\tau_K^\epsilon}B_{t,\tau_K^\epsilon}(\xi^\epsilon))\big]\cdot(1-2\epsilon C_2) + (\exp(-C_1-C_3) - \exp(C_3))\epsilon}{\bar{\mathbb{E}}_{t,y,i}\big[\exp(-\alpha_{\tau_K^\epsilon}B_{t,\tau_K^\epsilon}(\xi^\epsilon))\big]\cdot(1+2\epsilon C_2) + \exp(C_1+C_3)\epsilon} \\
&\ge -\frac{4C_2 + 4\exp(C_1+C_3)}{\alpha_t}\epsilon,
\end{aligned}$$

where the last inequality uses $\ln(1-x) \ge 1-2x$ for x small enough. Since ϵ was arbitrary this implies $\lim_{k\to\infty} p^k(t,y,i) \ge p_{t,T}(y,i)$. The fact that $p_{t,T}$ is a fixed point of $\mathcal{G}$ easily follows from the increasing property of $\mathcal{G}$: if $w^1 \ge w^2$ then $\mathcal{G}w^1 \ge \mathcal{G}w^2$. Since $p_{t,T}(y,i) \ge p^k(t,y,i)$ we have

$$\begin{aligned}
\mathcal{G}p_{t,T}(y,i) &\ge \lim_{k\to\infty}\mathcal{G}p^k(t,y,i) = \lim_{k\to\infty} p^{k+1}(t,y,i) = p_{t,T}(y,i) \\
&\ge \lim_{k\to\infty} p^k(t,y,i) = \lim_{k\to\infty}\mathcal{G}p^{k-1}(t,y,i) = \mathcal{G}p_{t,T}(y,i),
\end{aligned}$$

which implies that all inequalities must be equalities and therefore (3.12). The increasing property of $\mathcal{G}$ also implies that $p_{t,T}(\cdot,i)$ is the *smallest* fixed point of $\mathcal{G}$ bigger than $p^0(t,\cdot,i)$.

The bound on $p_{t,T}$ is immediate by

$$\begin{aligned}
p_{t,T}(y,i) &= -\frac{1}{\alpha_t}\ln\bar{\mathbb{E}}_{t,y,i}[\exp(-\alpha_T B_{t,T}(\xi^*))] \\
&\le -\frac{1}{\alpha_t}\ln\bar{\mathbb{E}}_{t,y,i}[\exp(-\alpha_T(T-t)\mathrm{e}^{r(T-t)}C_\psi)] = (T-t)C_\psi.
\end{aligned}$$

To establish the Lipschitz property of $p_{t,T}$ we recall that by Assumption 3, the flow $y \to Y_t^y(\omega)$ is locally Lipschitz, so that for any compact ball $B(y,r)$, and $t' \in [0,t]$ there is a constant $C(y)$ such that $Y_{t'}^{y_2} - Y_{t'}^{y_1} \le C(y)$ $\mathbb{P}$-a.s. for any $y_1, y_2 \in B(y,r)$. By the boundedness of ψ_i, we have the similar estimate

$|B_{t,T}(\xi;y_1) - B_{t,T}(\xi;y_2)| \le (T-t)e^{rT}C_\psi C(y) =: C_1$ where we explicitly show the dependence of the realized gains on the initial condition of $\{Y_t\}$.

Let ξ^1 be an optimal strategy for $p_{t,T}(y_1,i)$. We obtain

$$\begin{aligned}
p_{t,T}(y_1,i) - p_{t,T}(y_2,i) &\le p_{t,T}(y_1,i) - p_{t,T}(y_2,i;\xi^1) \\
&= -\frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,y_1,i}[\exp(-\alpha_T B_{t,T}(\xi^1))]}{\bar{\mathbb{E}}_{t,y_2,i}[\exp(-\alpha_T B_{t,T}(\xi^1))]} \\
&= -\frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,\cdot,i}[\exp(-\alpha_T(B_{t,T}(y_1,\xi^1) - B_{t,T}(y_2,\xi^1))\exp(-\alpha_T B_{t,T}(y_2,\xi^1))]}{\bar{\mathbb{E}}_{t,y_2,i}[\exp(-\alpha_T B_{t,T}(\xi^1))]} \\
&\le -\frac{1}{\alpha_t}\ln\frac{\bar{\mathbb{E}}_{t,\cdot,i}[\exp(-\alpha_T C_1|y_2-y_1|)\cdot\exp(-\alpha_T B_{t,T}(y_2,\xi^1))]}{\bar{\mathbb{E}}_{t,y_2,i}[\exp(-\alpha_T B_{t,T}(\xi^1))]} \\
&= \frac{\alpha_T}{\alpha_t}C_1|y_2-y_1|.
\end{aligned}$$

Repeating the same argument using ξ^2 establishes the opposite inequality and we conclude that $|p_{t,T}(y_1,i) - p_{t,T}(y_2,i)| \le e^{rT}C_1|y_2-y_1|$ uniformly in t and on compact subsets in y. □

References

L. Andersen and M. Broadie. A primal-dual simulation algorithm for pricing multi-dimensional American options. *Management Science*, 50(9):1222–1234, 2004.

V. Bally, G. Pagés, and J. Printems. A quantization tree method for pricing and hedging multidimensional American options. *Mathematical Finance*, 15(1):119–168, 2005.

K. Brekke and B. Øksendal. A verification theorem for combined stochastic control and impulse control. In *Stochastic analysis and related topics, VI (Geilo, 1996)*, pages 211–220. Birkhäuser, Boston, 1998.

M. Brennan and E. Schwartz. Evaluating natural resource investments. *Journal of Business*, 58:135–157, 1985.

R. Carmona and A. Danilova. Hedging financial instruments written on non-tradable indexes. Unpublished manuscript, 2003.

R. Carmona and M. Ludkovski. Pricing operational flexibility using optimal switching, 2005. Working paper.

R. Carmona, ed. *Volume on Indifference Pricing*. Princeton University Press, 2006.

P. Cheridito and M. Kupper. Time consistency of indifference prices and monetary utility functions, 2006. working paper.

P. Cheridito, F. Delbaen, and M. Kupper. Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, 11:57–106, 2006.

S. Dayanik and M. Egami. Solving impulse control problems via optimal stopping for one-dimensional diffusions. submitted, 2004.

A. Dixit. Entry and exit decisions under uncertainty. *Journal of Political Economy*, 97(3):620–638, June 1989.

A. Dixit and R. S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, 1994.

P. Dupuis and H. Wang. On the convergence from discrete to continuous time in an optimal stopping problem. *Annals of Applied Probability*, 15:1339–1366, 2005.

N. El Karoui. Les aspects probabilistes du contrôle stochastique. In *Ninth Saint Flour Probability Summer School—1979 (Saint Flour, 1979)*, volume 876 of *Lecture Notes in Math.*, pages 73–238. Springer, Berlin, 1981.

N. El Karoui and M. Jeanblanc. Optimization of consumption with labor income. *Finance and Stochastics*, 2:409–440, 1998.

A. Eydeland and K. Wolyniec. *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley&Sons, Hoboken, NJ, 2003.

E. Gobet, J. Lemor, and X. Warin. A regression-based Monte-Carlo method to solve backward stochastic differential equations. *The Annals of Applied Probability*, 15(3):2172–2002, 2005.

L. Györfi, M. Kohler, A. Krzyzak, and H. Walk. *A distribution-free theory of non-parametric regression*. Springer Series in Statistics, 2002.

S. Hamadène and M. Jeanblanc. On the starting and stopping problem: application in reversible investments. *Mathematics of Operations Research*, 32(1):182–192, 2007.

V. Henderson. Valuation of claims on nontraded assets using utility maximization. *Mathematical Finance*, 12(4):351–373, 2002.

V. Henderson. Explicit solutions to an optimal portfolio choice problem with stochastic income. *Journal of Economic Dynamics and Control*, 29(7):1237–1266, 2005.

V. Henderson and D. Hobson. Horizon-unbiased utility functions. *Stochastic Processes and their Applications*, 2007. in Press.

A. Ilhan, M. Jonsson, and R. Sircar. Portfolio optimization with derivatives and indifference pricing. In R. Carmona, editor, *Volume on Indifference Pricing*. Princeton University Press, 2006.

I. Karatzas and H. Wang. Utility maximization with discretionary stopping. *SIAM Journal on Control and Optimization*, 39(1):306–329, 2000.

S. Klöppel and M. Schweizer. Dynamic utility indifference valuation via convex risk measures. *Mathematical Finance*, 17(4):599–627, 2007.

H. J. Kushner and P. Dupuis. *Numerical methods for stochastic control problems in continuous time*, volume 24 of *Applications of Mathematics*. Springer-Verlag, New York, second edition, 2001.

J. Lepeltier and B. Marchal. Théorie générale du contrôle impulsionnel markovien. *SIAM Journal on Control and Optimization*, 22(4):645–665, 1984.

F. Longstaff and E. Schwartz. Valuing American options by simulations: a simple least squares approach. *Review of Financial Studies*, 14:113–148, 2001.

V. Ly Vath, M. Mnif, and H. Pham. A model of portfolio selection under liquidity risk and price impact. *Finance and Stochastics*, 11:51–90, 2007.

J. Miao and N. Wang. Investment, consumption and hedging under incomplete markets. *Journal of Financial Economics*, 2007. in Press.

M. Musiela and T. Zariphopoulou. A valuation algorithm for indifference prices in incomplete markets. *Finance and Stochastics*, 8(3):399–414, 2004.

B. Øksendal. *Stochastic differential equations*. Springer-Verlag, Berlin, fifth edition, 1998. An introduction with applications.

M. Owen and G. Žitković. Optimal investment with an unbounded random endowment and utility-based pricing methods. Technical report, University of Texas, 2007. Available at http://arxiv.org/abs/0706.0478.

A. Porchet, N. Touzi, and X. Warin. Valuation of power plants under optimal switching. Technical report, 2007. Available at http://www.cmap.polytechnique.fr/ touzi/ptw06.pdf.

E. Schwartz. The stochastic behavior of commodity prices:implications for valuation and hedging. *Journal of Finance*, LII(3):922–973, 1997.

S. Stoikov and T. Zariphopoulou. Optimal investments in the presence of unhedgeable risks and under CARA preferences. Technical report, U. Texas, 2005.

M. Tehranchi. Explicit solutions of some utility maximization problems in incomplete markets. *Stochastic Process. Appl.*, 114(1):109–125, 2004.

J. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, July 2001.

A. Yushkevich. Optimal switching problem for countable Markov chains: average reward criterion. *Math. Methods Oper. Res.*, 53(1):1–24, 2001.

T. Zariphopoulou. A solution approach to valuation with unhedgeable risks. *Finance and Stochastics*, 5(1): 61–82, 2001.

M. Zervos. A problem of sequential entry and exit decisions combined with discretionary stopping. *SIAM Journal on Control and Optimization*, 42(2):397–421, 2003.

Assistant Professor, Department of Mathematics, University of Michigan 530 Church St., Ann Arbor MI 48109

E-mail address: mludkov@umich.edu