Journal of Energy Markets **11**(4), 51–73
DOI: 10.21314/JEM.2018.183

Research Paper

Technical uncertainty in real options with learning

Ali Al-Aradi,[1] Álvaro Cartea[2,3] and Sebastian Jaimungal[1]

[1]Department of Statistical Sciences, University of Toronto, 100 St George Street, Toronto, Ontario M5S 3G3, Canada; emails: ali.al.aradi@utoronto.ca, sebastian.jaimungal@utoronto.ca

[2]Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter (550), Woodstock Road, Oxford OX2 6GG, UK; email: Alvaro.Cartea@maths.ox.ac.uk

[3]Oxford–Man Institute of Quantitative Finance, University of Oxford, Eagle House, Walton Well Road, Oxford OX2 6ED, UK

(Received June 25, 2017; revised February 23, 2018; accepted July 10, 2018)

ABSTRACT

We introduce a new approach for incorporating uncertainty in the decision to invest in a commodity reserve. An investment is an irreversible one-off capital expenditure, after which the investor receives a stream of cashflow from extracting the commodity and selling it on the spot market. The investor is exposed to price uncertainty as well as uncertainty in the amount of available resources in reserve (also known as "technical uncertainty"). They do, however, learn about the reserve levels over time and this is a key determinant in the decision to invest. To model the uncertainty surrounding the reserve levels and how the investor learns via estimates of the commodity in the reserve, we adopt a continuous-time Markov chain model; this allows us to value the option to invest in the reserve and to investigate the value of learning prior to investment.

Keywords: real option; investment under uncertainty; technical uncertainty; irreversibility; Markov chains.

Corresponding author: S. Jaimungal

Print ISSN 1756-3607 | Online ISSN 1756-3615

1 INTRODUCTION

What are the optimal market conditions for making an investment decision? This is an extensively studied problem in the academic literature and a key question at the heart of the valuation and execution of projects under uncertainty. Some investment projects are endowed with the option to delay decisions until market conditions are optimal. This option is valuable because it means that decisions may be taken when the potential gains stemming from them are at their maximum. The seminal work of McDonald and Siegel (1986) was the first to formalize the investment problem as a real option to invest in a project. In their paper, the value O_t of the option is calculated by comparing the difference between the value of investing now and the value of making the investment at a future time. Specifically, the value of the real option is written as

$$O_t = \sup_{\tau \in \mathcal{T}} \mathbb{E}[\mathrm{e}^{-\rho(\tau-t)}(V_\tau - I_\tau)_+], \tag{1.1}$$

where $\mathcal{T}$ is the set of admissible $\mathcal{F}$-stopping (exercise) times, ρ is the risk-adjusted discount rate, and V_t and I_t are the project value and (sunk) cost, respectively. The project value and cost are traditionally modeled using a geometric Brownian motion (GBM). The solution to this problem shows that the optimal investment strategy is to invest when the ratio V_t/I_t reaches a critical boundary B (the problem of optimal scrapping/divesting is similar, with the roles of V_t and I_t reversed). More recently, several authors have studied this problem with mean-reverting project values and costs (see, for example, Jaimungal *et al* 2013; Metcalf and Hasset 1995; Sarkar 2003).

In another influential paper, Brennan and Schwartz (1985) focus on the management of a mine (considering how the output rate should be controlled, the opening/ closing of the mine, abandonment and so on) rather than the optimal timing problem (see also Dixit 1989). They show how management decisions are modulated by output prices, which may be modeled as a GBM, assuming that costs are known.

These classical works do not take into account the uncertainty associated with reserve levels. To account for such "technical uncertainty", Pindyck (1980) develops a model in which the demand and reserve levels fluctuate continuously with increasing variance. Further, the optimal strategy is influenced by exploration and is introduced as a policy (ie, control) variable in two distinct ways. The first allows exploratory effort to affect the level of "knowledge", thus reducing the variance of the reserve fluctuations. The second assumes that reserves are discovered at a rate that depends on how much has been discovered in the past, the degree of effort currently being made to build on this progress and any exogenous noise.

More recent approaches to the investment timing problem with technical uncertainty include

- using Bayesian updating as in Armstrong *et al* (2004),
- modeling project costs via Markov chains as in Elliott *et al* (2009), and
- using proportionality to model learning as in Sadowsky (2005).

Cortázar *et al* (2001) describe a comprehensive approach to valuing several-stage exploration. As well as solving the timing problem, they also provide investment management decision rules pertaining to closure, opening and more (see also Brennan and Schwartz 1985). Other works incorporating real option techniques in the valuation of flexibility and investment decisions in the field of commodities and energy include the following:

- Himpler and Madlener (2014), who consider the optimal timing of wind farm repowering;
- Taschini and Urech (2010), who look at the option to switch fuels under different scenarios and fuel incentives;
- Fleten *et al* (2011), who examine the option to choose the capacity of an electricity interconnector between two locations; and
- Cartea and González-Pedraz (2012), who value an electricity interconnector by considering it as a stream of real options of the difference of prices in two locations.

Meanwhile, Cartea and Jaimungal (2017) study the effects of model uncertainty on irreversible investments. This work adds to the literature by incorporating both market and reserve uncertainty, while allowing the agent to learn about the status of reserves. Reserve uncertainty is represented by a Markov chain model with transition rates that decay as time flows forward to mimic the notion of learning. The model setup is developed in the context of oil exploration; however, it may be applied to other investment problems in commodities, such as mining for precious/base metals and natural gas fields. We value the irreversible option to invest in the exploration by developing a version of the Fourier space time stepping approach, as in Jackson *et al* (2008), Jaimungal and Surkov (2011) and Jaimungal and Surkov (2013) for equity, commodity and interest rate derivatives, respectively. For further studies on ambiguity aversion and model uncertainty in commodities and algorithmic trading, we refer the reader to Cartea *et al* (2016, 2017, 2018).

As long as estimates of the exploration costs and volumes are available, calibration of the model should be relatively simple. We demonstrate how the model can be used to assess whether exploration costs warrant the potential benefits to be gained from finding reserves and extracting them. Specifically, we show how to calculate the value of the option to delay investment and we discuss the agent's optimal investment threshold. This threshold, also referred to as the exercise boundary, depends on a number of variables and factors, including the agent's estimate of the volume of the reserves, the rate at which the agent learns about the volume of the reserves, the rate at which the agent extracts the commodity and the expiry of the option. We show that the value of the option to "wait and learn" is high at the beginning and gradually decreases as expiry of the option approaches and the agent is left with less and less time to learn.

We assume that the investment cost depends on the volume of the reserves. If the estimated volume is high (low), then the sunk cost of extracting the commodity will be similarly high (low). This has an effect on the optimal time for investment as well as the level of the spot price commodity required to justify incurring the sunk cost. For example, we show that when the volume estimate is low and the option to invest is far from the time of expiry, the agent sets a high investment threshold due to the effects of these two factors, thus making the option to delay investment valuable. First, a low volume requires a high commodity spot price to justify the investment. Second, as time to maturity increases the investor attaches a high value to waiting and learning the volume estimates of the reserve. However, as the option approaches expiry, the effects of these two factors become weaker. In particular, the value of learning falls because there is less time left for the investor to learn about the reserve estimates. Instead, the investment decision will merely be based on whether or not the investor's costs can be recovered; this depends, in turn, on the spot price of the commodity, the rate at which the commodity can be extracted and the uncertainty around it.

The remainder of this paper is organized as follows. In Section 2, we provide details of our modeling framework, explaining how we model both technical uncertainty and the uncertainty implicit in the underlying project. Moreover, we provide an approach to account for the agent's learning about the reserve environment through exploration. Section 3 provides an analysis of the Fourier space time stepping approach for valuing the early exercise features in the irreversible investment with learning. Section 4 shows how the model can be calibrated to estimates of the cost of exploration and the expected benefits of this exploration. Section 5 outlines various numerical experiments for demonstrating the efficacy of the approach as well as an analysis of the qualitative behavior of the model and its implications. Finally, we provide some concluding remarks and ideas for future lines of work.

2 MODEL ASSUMPTIONS

In this section, we provide models for the two sources of uncertainty that drive the value of the real option to explore and (irreversibly) invest in a project. The setting described here is tuned (to some extent) to oil exploration; however, it can be modified to deal with other activities, including mining, natural/shale gas and other natural explorations. Another extension to our setup is to account for the option to mothball exploration and/or extraction (once extraction has begun) among other managerial flexibilities that might arise in exploration and investment (see, for example, Dixit and Pindyck 1994; Jaimungal and Lawryshyn 2015; Kobari *et al* 2014; McDonald and Siegel 1985; Trigeorgis 1996; Tsekrekos 2010). We first describe how we model technical uncertainty and then describe how we model project value uncertainty. As usual, we assume a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geqslant 0}, \mathbb{P})$, where $\{\mathcal{F}_t\}_{t \geqslant 0}$ denotes the filtration (which will be described in more detail at a later stage). We also assume the existence of an equivalent martingale (or risk-neutral) measure, $\mathbb{Q}$.

2.1 Technical uncertainty

Let $V = (V_t)_{t \geqslant 0}$ denote the estimated reserve volume (level) process and ϑ be the true reserve volume. The true reserve volume ϑ is unknown to the investor and can be viewed as a random variable when conditioned on the information available to them at time t.[1] It will, however, be revealed as $t \to \infty$. For simplicity, we assume that the possible reserve levels (as well as their estimates) take on values from a finite set of possible reserve volumes. We model the estimated reserve level as a continuous-time (inhomogeneous) Markov chain in view of the fact that reserve estimates are being constantly updated as new information is obtained via exploration. Moreover, to capture the tendency of the accuracy of estimates to improve over time (as more information becomes available), we assume that the transition rate between volume estimate states will decrease as time progresses. Further, we assume that $\lim_{t \to \infty} V_t = \vartheta$ almost surely to reflect the feature mentioned earlier; that is, the true reserve level is revealed to the investor with the passage of time.[2]

More specifically, the estimated reserve volume V_t is modulated by a finite state, continuous-time, inhomogeneous Markov chain Z_t, taking values in $\{1, \dots, m\}$ via

$$V_t = v^{(Z_t)}, \tag{2.1}$$

where the constants

$$\{v^{(1)}, \dots, v^{(m)}\} \in \mathbb{R}_+^m \tag{2.2}$$

[1] For which we need to assume that $\vartheta < \infty$ almost surely.

[2] We could in principle derive such a model by writing $V_t = \mathbb{E}[\vartheta \mid \mathcal{G}_t]$, where $\mathcal{G} = (\mathcal{G}_t)_{t \geqslant 0}$ denotes the filtration used to generate information about the true reserve. Under specific modeling assumptions, V can be cast into a Markov chain representation. We opt not to delve into these details (this would detract from the simplicity of the approach we are proposing) and instead model V directly.

are the possible reserve volumes. The generator matrix of the Markov chain Z_t is denoted by $\boldsymbol{G}_t$ and assumed to be of the form

$$\boldsymbol{G}_t = h_t \boldsymbol{A}, \tag{2.3}$$

where h_t is a deterministic, nonnegative decreasing function of time such that $h_t \xrightarrow{t\to\infty} 0$ and $\int_0^\infty h_u \mathrm{d}u < \infty$. Further, $\boldsymbol{A}$ is a constant $m \times m$ matrix with $\sum_{j=1}^{n} A_{ij} = 0$ and $A_{ij} > 0$ for $i \neq j$. The states of the Markov chain correspond with various possible estimates for the reserve level, thus capturing the uncertainty of these estimates.

The function h_t captures how the agent learns about the volume (or quantity) of the commodity in the reserves. A decreasing h implies that the transition rates are also decreasing; hence the probability of any change in the estimated volume decreases with time and the estimate becomes more accurate. Optimal policies for the irreversible investment that could be explored as well as the value of the project following extraction of the commodity naturally depend on the observed estimate of the reserves. Section 4 discusses the forms of h_t and $\boldsymbol{A}$ in detail as well as the process of their calibration to the data.

2.2 Market uncertainty

The second source of uncertainty stems from the spot price of the commodity. We denote this spot price as $S = (S_t)_{t\geqslant 0}$. It may be modeled as the exponential of an Ornstein–Uhlenbeck (OU) process:

$$S_t = \exp\{\theta + X_t\}. \tag{2.4a}$$

In (2.4a), the OU process $X = (X_t)_{t\geqslant 0}$ satisfies the stochastic differential equation (SDE)

$$\mathrm{d}X_t = -\kappa X_t \,\mathrm{d}t + \sigma \,\mathrm{d}W_t, \tag{2.4b}$$

where $W = (W_t)_{t\geqslant 0}$ is a standard Brownian motion, $\kappa > 0$ is the rate of mean reversion, θ is the (log) level of the mean reversion and σ is the (log) volatility of the spot price. Such models of commodity spot prices are widely used in the literature (see, for example, Cartea and Figueroa 2005; Cartea and Williams 2008; Coulon *et al* 2013; Kiesel *et al* 2009; Weron 2007).

Now that we have specified the model for the stock of the commodity in the reserve and its market price, we need one final ingredient, namely, the market value of the commodity in the reserve. We denote this value by the process $P = (P_t)_{t\geqslant 0}$ and show how to calculate it in steps.

Suppose that an investment is made at time t and is followed by extraction of the commodity $\varepsilon \geqslant 0$ later, and that the extraction process continues until the random (stopping) time $\tau = t + \varepsilon + \Delta$. $\Delta > 0$ represents the amount of time required to complete the extraction process. Since the quantity of the commodity in the reserves is unknown to the investor along with the time to depletion (and the extraction duration), the stopping time is random. We also note that engineering and physical limitations may prevent the total amount of the commodity stored in the reserve from being extracted; instead, only $\gamma\vartheta$ is extractable, $0 < \gamma < 1$.[3]

The random time to depletion Δ is related to the unknown total reserve volume ϑ via the rate of extraction. In particular, we have

$$\int_{t+\varepsilon}^{t+\varepsilon+\Delta} g(u)\,\mathrm{d}u = \gamma\vartheta, \tag{2.5}$$

where $g(u)$ denotes the rate of extraction at time u. Once extraction begins, we assume that the commodity will be extracted at the rate

$$g(u) = \alpha \mathrm{e}^{-\beta(u-(t+\varepsilon))}, \qquad u \in [t+\varepsilon, t+\varepsilon+\Delta], \tag{2.6}$$

where $\alpha \geqslant 0$ and $\beta \geqslant 0$. Figure 1 presents a stylized picture of the exponential extraction rate (2.6). Under the specific extraction rate model given in (2.6), the time to depletion can be written as

$$\Delta = -\frac{1}{\beta}\log\left(1 - \frac{\beta}{\alpha}\gamma\vartheta\right), \tag{2.7}$$

which is a random variable because ϑ is known only as $t \to \infty$.

The value of the reserve when extraction begins is determined by a number of factors, including the price of the commodity, the state of the Markov chain linked to reserve uncertainty and the random time to exhaustion of the reserve. Specifically, the discounted value of the cashflow generated from extracting the commodity at the rate $g(u)$ and selling it at the spot price S_u is given by

$$\mathrm{DCF}_t = \int_{t+\varepsilon}^{t+\varepsilon+\Delta} \mathrm{e}^{-\rho(u-t)}(S_u - c)g(u)\,\mathrm{d}u, \tag{2.8}$$

where ρ is the agent's discount factor, given the level of risk they bear with the project, and c is a running cost incurred by the investor as long as the extraction operation continues.

To compute the market value of the commodity in the reserve, which we denote by P_t, we calculate the expected discounted value of the cashflows attained from selling the extracted commodity. Namely, we insert into (2.8) the spot price of the

[3] This is particularly relevant to the example of stored natural gas, where there is always a residual that cannot be extracted from storage.

FIGURE 1 Once the agent invests in the reserve at time t, extraction begins at time $t + \varepsilon$ and continues until γ% of the actual (unknown) reserve volume ϑ is extracted; this occurs at the random time $t + \varepsilon + \Delta$.

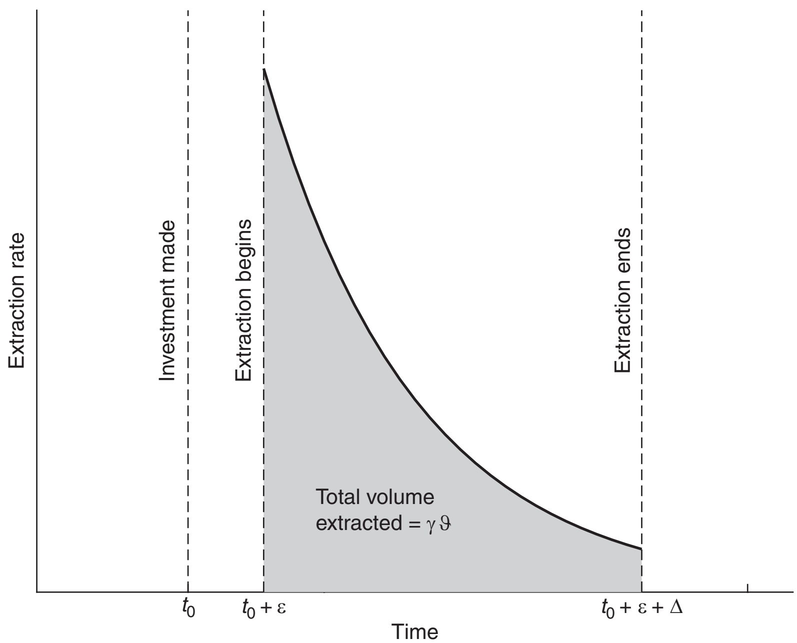

commodity (2.4), the time to extraction completion given in (2.7) and the extraction rate (2.6). Finally, we compute the expectation of DCF_t to obtain

$$P_t = \mathbb{E}[\mathrm{DCF}_t \mid \mathcal{F}_t] = \mathbb{E}\biggl[\int_{t+\varepsilon}^{t+\varepsilon+\Delta} \mathrm{e}^{-\rho(u-t)}(F_t(u) - c)g(u)\,\mathrm{d}u \biggm| \mathcal{F}_t\biggr], \qquad (2.9)$$

where $\mathcal{F}_t = \sigma((S_u, V_u)_{u \in [0,t]})$ is the natural filtration generated by both S and V (or, equivalently, the Markov chain Z introduced earlier), the expectation is taken under the risk-neutral measure $\mathbb{Q}$, and

$$\begin{aligned} F_t(u) &= \mathbb{E}[S_u \mid \mathcal{F}_t] \\ &= \exp\biggl\{\theta + \mathrm{e}^{-\kappa(u-t)}x + \frac{\sigma^2}{4\kappa}(1 - \mathrm{e}^{-2\kappa(u-t)})\biggr\} \end{aligned}$$

is the forward price of the underlying asset.

Thus, the expectation in (2.9) concerns the random time to depletion Δ, which further depends on the reserve volume ϑ. To calculate this expectation, we must obtain the conditional distribution of the unknown reserve level ϑ, that is, we require $\mathbb{P}(\vartheta = v^{(j)} \mid \mathcal{F}_t)$. It follows from our assumption that $\vartheta = \lim_{t\to\infty} V_t$ almost surely that this, in turn, requires us to determine the conditional distribution of the limit of the underlying Markov chain at time t. Specifically,

$$\begin{aligned}\mathbb{P}(\vartheta = v^{(j)} \mid \mathcal{F}_t) &= \mathbb{P}\Big(\lim_{t\to\infty} V_t = v^{(j)} \,\Big|\, \mathcal{F}_t\Big)\\ &= \mathbb{P}\Big(\lim_{t\to\infty} Z_t = j \,\Big|\, Z_t = i\Big)\\ &= [\mathrm{e}^{H_t \boldsymbol{A}}]_{ij},\end{aligned}$$

where $H_t = \int_t^\infty h_u \,\mathrm{d}u$. The notation $[\cdot]_{ij}$ denotes the ij element of the matrix in the square brackets, while the matrix $\boldsymbol{A}$ is defined in (2.3). Therefore, the final expression for the expected discounted value of the reserve is

$$\begin{aligned}P_t &:= p^{(Z_t)}(t, X_t)\\ &= \sum_{j=1}^{m} [\mathrm{e}^{H_t \boldsymbol{A}}]_{Z_t, j} \int_{t+\varepsilon}^{t+\varepsilon-(1/\beta)\log(1-(\beta/\alpha)\gamma v^{(j)})} \mathrm{e}^{-\rho(u-t)}(F_t(u) - c)g(u)\,\mathrm{d}u.\end{aligned} \tag{2.10}$$

This expression has two sources of uncertainty: the first stems from the spot price of the commodity, through the OU process X_t, and the second stems from the estimate of the reserve volume, through the state of the Markov chain Z_t. If we assume the exponentially decaying extraction rate model given in (2.6), it is possible to write the integral that appears in the right-hand side of (2.10) in terms of special functions. Such a rewrite, however, does not add much in terms of clarity, so we opt to keep the integral as shown above.

3 REAL OPTION VALUATION

Having established our model for the value of the reserve, we turn our attention to the cost of exploiting the reserves of the commodity as well as the value of having the flexibility to decide when to make an investment. The cost of investing in the reserve is irreversible and denoted by $I^{(k)}$, where k is the regime of the reserve volume estimate. Here, we assume that the cost $I^{(k)}$ is linked with the volume estimate since extracting a large reserve will likely require a larger investment up front than would be required to extract the commodity from a smaller reserve. We assume that the investment cost is

$$I^{(k)} = c_0 + c_1 v^{(k)}, \tag{3.1}$$

where $c_0 \geqslant 0$ is a fixed cost, $c_1 \geqslant 0$, and $v^{(k)}$ are the possible reserve volumes (as we recall from (2.2)). Note that the cost of the investment depends on the reserve estimate. This reflects the fact that the investor will base the size of their extraction operation on their estimate of the reserve amount; that is, the larger they believe the reserve to be, the bigger the operation required. The running cost c ensures that the investor's running costs remain tied to the true reserve amount.

We denote the value of the option by $L_t = \ell^{(Z_t)}(t, X_t)$, where the collection of functions $\ell^{(1)}(t,x), \dots, \ell^{(m)}(t,x)$ represents the value of the real option conditional on the state $Z_t = 1, \dots, m$ (indexed by the superscript) and $X_t = x$. The agent must make their decision by time T; after this point, they lose their option to make an investment and exploit the reserve. Standard theory implies that the value of the option to irreversibly invest in the reserve and begin extraction is given by the optimal stopping problem:

$$L_t = \sup_{\tau \in \mathcal{T}} \mathbb{E}[\mathrm{e}^{-\rho\tau} \max(P_\tau - I^{(Z_\tau)}, 0) \mid \mathcal{F}_t] \tag{3.2a}$$

$$= \sup_{\tau \in \mathcal{T}} \mathbb{E}[\mathrm{e}^{-\rho\tau} \max(P_\tau - I^{(Z_\tau)}, 0) \mid Z_t, X_t]. \tag{3.2b}$$

Here, $\mathcal{T}$ denotes the set of admissible stopping times, taken to represent the finite collection of $\mathcal{F}$-stopping times restricted to $t_i = i\Delta t$, $i = 0, \dots, N$, with $t_N \leqslant T$. In other words, the agent is restricted to making their investment decision on days t_i. In the interim time, the agent can acquire further information to improve their volume reserve estimates.

For notational convenience, we define the deflated value process as

$$\bar{\ell}^{(j)}(t,x) := \mathrm{e}^{-\rho t} \ell^{(j)}(t,x) \tag{3.3}$$

and observe that, in between the investment dates, the deflated value processes $\bar{\ell}^{(j)}(t,x)$ for $j = 1, \dots, m$ are martingales. Since there is no opportunity to exercise the option in between the investment dates, $\bar{\ell}^{(j)}(t,x)$ is the same as a European claim, with a payoff that is equal to its value at the next exercise date. Thus,

$$\bar{\ell}^{(j)}(t_i, x) = \max\Big(\lim_{t \downarrow t_i} \bar{\ell}^{(j)}(t,x); \mathrm{e}^{-\rho t_i}(p^{(j)}(t,x) - I^{(j)})\Big), \tag{3.4}$$

where $p^{(j)}(t,x)$ is as in (2.10) and $j = 1, \dots, m$ represents the state of the regime as before.

Finally, in the interval $t \in (t_i, t_{i+1}]$ the processes $\bar{\ell}^{(j)}(t,x)$ satisfy the coupled system of partial differential equations (PDEs):

$$(\partial_t + \mathscr{L})\bar{\ell}^{(j)}(t,x) + h_t \sum_{j=1}^{m} A_{jk} \bar{\ell}^{(k)}(t,x) = 0, \quad t \in (t_i, t_{i+1}], \tag{3.5}$$

where $\mathscr{L} = -\kappa x \partial_x + \frac{1}{2}\sigma^2 \partial_{xx}$ is the infinitesimal generator of the process X_t.

The maximization in (3.4) represents the agent's option to either hold on to the investment option at time t_i or to invest immediately. If the second argument attains the maximum, then the agent exercises their option to invest in the reserve at a cost of $I^{(j)}$; they then receive the expected discounted value of the cashflow $p^{(j)}(t,x)$, which results from extracting and selling the commodity on the spot market. This investment decision is tied to the reserve volume estimate through the regime j. Different regimes j result in different exercise policies. We explore the relationship between the two in the next section.

Motivated by the work of Jaimungal and Surkov (2011), who study options on multiple commodities driven by Lévy processes, we solve the system of PDEs (3.5) recursively by employing a Fourier transform of $\bar{\ell}^{(j)}(t,x)$ with respect to x, denoted here as $\tilde{\ell}^{(j)}(t,\omega)$. Specifically, we write

$$\tilde{\ell}^{(j)}(t,\omega) = \int_{-\infty}^{\infty} \mathrm{e}^{-\mathrm{i}\omega x}\bar{\ell}^{(j)}(t,x)\,\mathrm{d}x \quad \text{and} \quad \bar{\ell}^{(j)}(t,x) = \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{i}\omega x}\tilde{\ell}^{(j)}(t,\omega)\frac{\mathrm{d}\omega}{2\pi}, \tag{3.6}$$

where $\mathrm{i} = \sqrt{-1}$. Applying the Fourier transform to (3.5), we obtain a new PDE (without the parabolic term) that depends on the state variable ω rather than the state variable x, that is,

$$[\partial_t + (\kappa - \tfrac{1}{2}\sigma^2\omega^2) + \kappa\omega\partial_\omega]\tilde{\ell}^{(j)}(t,\omega) + h_t\sum_{j=1}^{m} A_{jk}\tilde{\ell}^{(k)}(t,\omega) = 0. \tag{3.7}$$

Within the interval $(t_k, t_{k+1}]$, we introduce a moving coordinate system and write $\hat{\ell}^{(j)}(t,\omega) = \tilde{\ell}^{(j)}(t, \mathrm{e}^{-\kappa(t_{k+1}-t)}\omega)$; thus, removing the derivative in ω, we find that the functions $\hat{\ell}^{(j)}$ satisfy the coupled linear system of ordinary differential equations (ODEs):

$$\partial_t\hat{\ell}^{(j)}(t,\omega)+(\kappa-\tfrac{1}{2}\sigma^2\omega^2\mathrm{e}^{-2\kappa(t_{k+1}-t)})\hat{\ell}^{(j)}(t,\omega)+h_t\sum_{j=1}^{m} A_{jk}\hat{\ell}^{(k)}(t,\omega) = 0. \tag{3.8}$$

The coupled system of ODEs above can be recast as independent ODEs by writing $\boldsymbol{A} = \boldsymbol{U}\boldsymbol{D}\boldsymbol{U}^{-1}$, where $\boldsymbol{U}$ is the matrix of the eigenvectors of $\boldsymbol{A}$, while $\boldsymbol{D}$ is the diagonal matrix of the eigenvalues of $\boldsymbol{A}$. In vector form, this results in

$$\partial_t(\boldsymbol{U}^{-1}\hat{\boldsymbol{\ell}}(t,\omega)) + (\psi(\omega\mathrm{e}^{-\kappa(t_{k+1}-t)})\mathbb{I} + h_t\boldsymbol{D})\boldsymbol{U}^{-1}\hat{\boldsymbol{\ell}}(t,\omega) = \boldsymbol{0}, \tag{3.9}$$

where $\hat{\boldsymbol{\ell}}(t,\omega) = (\hat{\ell}^{(1)}(t,\omega),\dots,\hat{\ell}^{(n)}(t,\omega))'$, $\psi(\omega) = \kappa - \frac{1}{2}\sigma^2\omega^2$ and $\mathbb{I}$ is the $n\times n$ identity matrix. These uncoupled ODEs have the solution

$$\begin{aligned}\boldsymbol{U}^{-1}\hat{\boldsymbol{\ell}}(t_k^+,\omega) \\ &= \exp\left\{\int_{t_k}^{t_{k+1}} \psi(\omega\mathrm{e}^{-\kappa(t_{k+1}-s)})\,\mathrm{d}s\mathbb{I} + \int_{t_k}^{t_{k+1}} h_s\,\mathrm{d}s\boldsymbol{D}\right\}\boldsymbol{U}^{-1}\hat{\boldsymbol{\ell}}(t_{k+1},\omega),\end{aligned} \tag{3.10}$$

where $\hat{\boldsymbol{\ell}}(t_{k_+},\omega) = \lim_{t\downarrow t_k}\hat{\boldsymbol{\ell}}(t,\omega)$.

Next, we left-multiply by $\boldsymbol{U}$ to obtain

$$\hat{\boldsymbol{\ell}}(t_k^+,\omega) = \exp\left\{\int_{t_k}^{t_{k+1}} \psi(\omega \mathrm{e}^{-\kappa(t_{k+1}-s)})\,\mathrm{d}s\right\} \exp\left\{\int_{t_k}^{t_{k+1}} h_s\,\mathrm{d}s\boldsymbol{A}\right\}\hat{\boldsymbol{\ell}}(t_{k+1},\omega), \tag{3.11}$$

and the Fourier transform of the deflated value of the option to irreversibly invest is

$$\begin{aligned}\tilde{\boldsymbol{\ell}}(t_k^+,\omega) = \exp&\left\{\int_0^{t_{k+1}-t_k} \psi(\omega \mathrm{e}^{\kappa s})\,\mathrm{d}s\right\}\\ &\times \exp\left\{\int_{t_k}^{t_{k+1}} h_s\,\mathrm{d}s\boldsymbol{A}\right\}\tilde{\boldsymbol{\ell}}(t_{k+1},\omega \mathrm{e}^{\kappa(t_{k+1}-t_k)}).\end{aligned} \tag{3.12}$$

This result is interesting in a few respects. First, the role of the mean reversion decouples from the Markov chain driving the volume estimates. Second, the value at time t_k^+ and frequency ω depends on the value at time t_{k+1} and frequency $\omega \mathrm{e}^{\kappa(t_{k+1}-t_k)}$. An extrapolation of the frequency space is, therefore, required as the algorithm used to calculate the option value steps backward in time. Once we discretize the state space, such extrapolations can lead to inaccurate results since the edges of the state space constitute a most important contributing factor in calculating the extrapolated values. Instead, we make use of the inverse relationship between frequencies and real space in Fourier transforms,

$$\int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{i}\omega x} g(x/a)\,\mathrm{d}x = \int_{-\infty}^{\infty} \mathrm{e}^{\mathrm{i}(a\omega)x} g(x)\frac{\mathrm{d}x}{a} = \frac{1}{a}\tilde{g}(a\omega),$$

to write

$$\tilde{\boldsymbol{\ell}}(t_k^+,\omega) = \exp\left\{\int_0^{t_{k+1}-t_k} \psi(\omega \mathrm{e}^{\kappa s})\,\mathrm{d}s\right\} \exp\left\{\int_{t_k}^{t_{k+1}} h_s\,\mathrm{d}s\boldsymbol{A}\right\}\tilde{\breve{\boldsymbol{\ell}}}(t_{k+1},\omega),$$

where $\breve{\boldsymbol{\ell}}(t_{k+1},x) = \boldsymbol{\ell}(t_{k+1},x\mathrm{e}^{-\kappa(t_{k+1}-t_k)})$ and $\tilde{\breve{\boldsymbol{\ell}}}$ denotes the Fourier transform of $\breve{\boldsymbol{\ell}}$. Thus, the value of the right-limit of the real option to invest at t_k^+ is determined in terms of interpolation in x (rather than extrapolation in ω). In Algorithm 1, we summarize our approach for valuing the real option to irreversibly invest in the reserve.

4 CALIBRATION AND LEARNING

Armed with the model presented in Section 2 and the valuation procedure developed in Section 3, we discuss in detail the calibration procedure for the Markov chain parameters and the investor's learning function h_t given by (2.3). The general idea is to use the investor's prior information regarding the estimate to define the possible reserve estimates (represented by the states of the Markov chain) as well as

ALGORITHM 1. Algorithm for computing the value of the option to irreversibly invest in the reserve.

---

(1) Set the real and frequency space grids:

$$x = [-\bar{x}\colon \Delta x\colon \bar{x}], \quad \breve{x} = \mathrm{e}^{-\kappa \Delta t} x \quad \text{and} \quad \omega = [0\colon \Delta\omega\colon \bar{\omega}].$$

(2) Place terminal conditions: $\boldsymbol{\ell}(t_n, x) = \mathrm{e}^{-\rho t_n}(\boldsymbol{P}(t_n, x) - \boldsymbol{I})_+$.

(3) Set $k = n$.

(4) Step backward from t_{k+1} to t_k:

 (a) $\breve{\boldsymbol{\ell}}_{t_{k+1}} = \mathrm{interp}(x, \boldsymbol{\ell}_{t_{k+1}}, \breve{x})$;

 (b) $\tilde{\breve{\boldsymbol{\ell}}}_{t_{k+1}}(\omega) = F[\breve{\boldsymbol{\ell}}_{t_{k+1}}(x)]$;

 (c) $\breve{\boldsymbol{\ell}}_{t_k^+}(x) = F^{-1}\biggl[\exp\biggl\{\int_0^{t_{k+1}-t_k} \psi(\omega \mathrm{e}^{\kappa s})\,\mathrm{d}s\biggr\}\exp\biggl\{\int_{t_k}^{t_{k+1}} h_s\,\mathrm{d}s\boldsymbol{A}\biggr\}\tilde{\breve{\boldsymbol{\ell}}}_{t_{k+1}}(\omega)\biggr]$;

 (d) $\tilde{\boldsymbol{\ell}}_{t_k} = \max(\tilde{\boldsymbol{\ell}}_{t_k^+}(x); \mathrm{e}^{-\rho t_k}(\boldsymbol{P}(t_k, x) - \boldsymbol{I})_+)$.

(5) Set $k \to k-1$, if $k \geqslant 0$ go to step (4).

---

the transition rates for the base generator matrix $\boldsymbol{A}$. The next step is to calibrate the learning function using information about how much the reserve estimate variance can be reduced by some future date.

4.1 Calibration of the Markov chain parameters

At time $t = 0$, the agent has obtained an estimate of the true reserve volume, denoted by μ, and the volatility of the estimate, denoted by σ_0. Thus,

$$\mu = \mathbb{E}[\vartheta \mid \mathcal{F}_0] \quad \text{and} \quad \sigma_0^2 = \mathbb{V}[\vartheta \mid \mathcal{F}_0]. \tag{4.1}$$

The first step in the calibration procedure is to select the states of the reserve volume conditional on the Markov chain state, that is, to select $v^{(1)}, \dots, v^{(m)}$. Since the $t = 0$ estimate of the reserve volume represents an unbiased estimator of the reserves, the Markov chain should be symmetric around the initial estimate of the reserve volume. To ensure this symmetry, we assume the following:

(i) the cardinality of the states of the Markov chain Z_t is odd, ie, $m = 2L + 1$ for some positive integer L;

(ii) $v^{(L+1)} = \mu$, where μ is the initial estimate of the reserve volume;

(iii) $v^{(k)}$ is increasing in k; and

(iv) $v^{(L+1+i)} - \mu = \mu - v^{(L+1-i)}$ for all $i = 1, \dots, L$.

We further assume the agent's estimate of the volume to be normally distributed (this assumption can be modified to fit with any distribution considered by the agent as representative of their prior knowledge):

$$\vartheta \mid \mathcal{F}_0 \sim \mathcal{N}(\mu, \sigma_0^2). \tag{4.2}$$

Therefore, we choose $v^{(k)}$ to be equally spaced over n standard deviations of the normal random variable support, that is, we select

$$v^{(k)} = \mu - n\sigma_0 + (k-1)\frac{2n\sigma_0}{m-1} \quad \text{for all } k = 1, \dots, m. \tag{4.3}$$

Placing symmetry on the states of the reserve volume estimator is not sufficient to ensure symmetry in its distribution. We also require symmetry in the base generator rate matrix $\boldsymbol{A}$ and assume that

$$\boldsymbol{A}^{\boldsymbol{\lambda}}_{1,1} = -\lambda_1, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{1,2} = \lambda_1, \tag{4.4a}$$

$$\boldsymbol{A}^{\boldsymbol{\lambda}}_{i,i-1} = \lambda_i, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{i,i} = -2\lambda_i, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{i,i+1} = \lambda_i, \quad i = 2, \dots, L, \tag{4.4b}$$

$$\boldsymbol{A}^{\boldsymbol{\lambda}}_{L+1,L} = \lambda_{L+1}, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{L+1,L+1} = -2\lambda_{L+1}, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{L+1,L+2} = \lambda_{L+1}, \tag{4.4c}$$

$$\boldsymbol{A}^{\boldsymbol{\lambda}}_{m-i,m-i-1} = \lambda_i, \;\; \boldsymbol{A}^{\boldsymbol{\lambda}}_{m-i,m-i} = -2\lambda_i, \;\; \boldsymbol{A}^{\boldsymbol{\lambda}}_{m-i,m-i+1} = \lambda_i, \;\; i = 1, \dots, L-1, \tag{4.4d}$$

$$\boldsymbol{A}^{\boldsymbol{\lambda}}_{m,m-1} = \lambda_1, \quad \boldsymbol{A}^{\boldsymbol{\lambda}}_{m,m} = -\lambda_1. \tag{4.4e}$$

for some set of $\boldsymbol{\lambda} = \{\lambda_1, \lambda_2, \dots, \lambda_{L+1}\}$, where $\lambda_1, \lambda_2, \dots, \lambda_{L+1} > 0$. Here, λ_1, $\{\lambda_2, \dots, \lambda_L\}$ and λ_{L+1} determine the transition rates out of the edge, interior and mid-states, respectively. The form of $\boldsymbol{A}^{\boldsymbol{\lambda}}$ ensures that transitions only occur between neighboring states; it also guarantees symmetry in the transition rates of states that are on either side of the mean estimate, ie, states $m+1$ and $m-1$ have the same transition rate, as do states $m+2$ and $m-2$, etc. The parameters $\boldsymbol{\lambda}$ are calibrated so that the invariant distribution of ϑ without learning coincides with a discrete approximation of a normal random variable with mean μ and variance σ_0^2. This ensures that the Markov chain generates an invariant distribution that is equal to a discrete approximation of the original estimate of the reserve volume distribution. Formally, let $\boldsymbol{P}^{\boldsymbol{\lambda}} = \mathrm{e}^{\boldsymbol{A}^{\boldsymbol{\lambda}}}$ denote the transition probability after one unit of time, and let $\boldsymbol{\pi}^{\boldsymbol{\lambda}}$ denote the invariant distribution of $\boldsymbol{P}$; that is, $\boldsymbol{\pi}^{\boldsymbol{\lambda}}$ solves the eigenproblem $\boldsymbol{P}\boldsymbol{\pi}^{\boldsymbol{\lambda}} = \boldsymbol{\pi}^{\boldsymbol{\lambda}}$. Then, we choose $\boldsymbol{\lambda}$ such that

$$\boldsymbol{\pi}^{\boldsymbol{\lambda}}_1 = \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(1)} + v^{(2)})) - \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(1)} + [v^{(1)} - (v^{(2)} - v^{(1)})])), \tag{4.5a}$$

$$\boldsymbol{\pi}^{\boldsymbol{\lambda}}_2 = \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(i+1)} + v^{(i)})) - \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(i)} + v^{(i-1)})), \quad i = 2, \dots, 2L, \tag{4.5b}$$

$$\boldsymbol{\pi}^{\boldsymbol{\lambda}}_m = \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(m)} + [v^{(m)} + (v^{(m)} - v^{(m-1)})])) - \Phi_{\mu,\sigma_0}(\tfrac{1}{2}(v^{(m)} + v^{(m-1)})), \tag{4.5c}$$

where $\Phi_{\mu,\sigma_0}(\cdot)$ denotes the normal cumulative density of a normal with mean μ and variance σ_0^2. Note that Φ_{μ,σ_0} is evaluated at the midpoint between the possible reserve amounts, and that we extrapolate linearly to obtain equations for the edge states.

4.2 Calibration of the learning function

As time passes, the agent gathers higher quality information (and more of it) about the volumes of the commodity in the reserve. Thus, the variance of the estimated volume in the reserve is expected to decrease from σ_0^2 to $\sigma_{T'}^2 < \sigma_0^2$ by some fixed time $T' < T$. Specifically,[4]

$$\sigma_{T'}^2 = \mathbb{V}[\vartheta \mid \mathcal{F}_{T'}]. \tag{4.6}$$

This parameter and the starting reserve estimate variance σ_0^2 constitute the main determinants of the learning function. For parsimony, we assume that the agent's learning function is of the form

$$h_t = a\mathrm{e}^{-bt} \quad \text{for some } a, b > 0,$$

where the parameter a represents the initial transition rates (ie, $h_0 = a$) between the states of the Markov chain and hence reflects the uncertainty of the initial estimates of the reserves. Learning parameter b represents the rate at which the agent learns; the larger (smaller) this value, the quicker (slower) the learning process; this is because large (small) values of b make the transition rates decay faster (slower) through time. Therefore, the reserve estimates become stable more quickly (slowly). Recall that the learning rate function plays a key role in the generator matrix of the Markov chain (see (2.3)) in that it captures how the agent learns the volume of the commodity being held in the reserve.

The parameters in the learning rate function h are calibrated to obey constraints (4.1) and (4.6). Due to the symmetry in the base transition rate $\boldsymbol{A}$ and in the reserve volume states $v^{(k)}$, we automatically satisfy the mean constraints

$$\mathbb{E}[\vartheta \mid V_0 = \mu] = \mu \quad \text{and} \quad \mathbb{E}[\vartheta \mid V_{T'} = \mu] = \mu. \tag{4.7}$$

To satisfy the variance constraints, we require the transition probabilities of the Markov chain from an arbitrary state at time t to its infinite horizon state, which we denote as $p_{t,ij} := \mathbb{P}(\lim_{t\to\infty} Z_t = j \mid Z_t = i)$, to be

$$p_{t,ij} = \left[\exp\left\{\left(\int_t^\infty h_u \,\mathrm{d}u\right)\boldsymbol{A}\right\}\right]_{ij} = \left[\exp\left\{\frac{a}{b}\mathrm{e}^{-bt}\boldsymbol{A}\right\}\right]_{ij}. \tag{4.8}$$

[4] In principle, we could develop a model that calibrates to a sequence of times and variances; however, the one-step reduction is enough to illustrate the essential ideas in this reduction.

Note that the right-hand side of the equation above is a matrix exponential; as before, the notation $[\cdot]_{ij}$ denotes the ij element of the matrix in the square brackets. Now we must solve a coupled system of two nonlinear equations for the parameters a and b via

$$\sigma_0^2 = \mathbb{V}[\vartheta \mid V_0 = \mu] = \sum_{k=1}^{2L+1} v^{(k)}(p_{0,Lk} - \mu)^2, \tag{4.9a}$$

$$\sigma_{T'}^2 = \mathbb{V}[\vartheta \mid V_{T'} = \mu] = \sum_{k=1}^{2L+1} v^{(k)}(p_{T',Lk} - \mu)^2, \tag{4.9b}$$

where $\mathbb{V}[\cdot \mid \cdot]$ is the variance operator of the first argument conditional on the event represented by the second argument.

All technical uncertainty model parameters must now be calibrated to the distributional properties of the initial reserve volume estimates as well as the reduction in variance that comes as a result of the agent's learning.

The value of the irreversible option to invest in the reserve with learning can now be valued using the approach outlined in Section 3, a summary of which is presented in Algorithm 1. The value of exploration (a process that improves the variance of the estimators of the reserve volume) can be assessed by considering the option value when there is no learning and comparing it with the option value with learning.

5 NUMERICAL RESULTS

In this section, we investigate the optimal exercise policies that may be adopted by the agent and assess the value of their learning. We use the following parameters and modeling choices throughout.

Reserve volume. Initial reserve estimate $\mu = 10^9$; initial reserve estimate variance $\sigma_0^2 = 3 \times 10^8$.

Investment costs. Fixed cost parameter $c_0 = 10^8$; variable cost parameter $c_1 = 3 \times 10^6$. This implies an investment cost of 1.12×10^8 when the reserve estimate is at its lowest and 1.48×10^8 when it is at its highest.

Expiry of option. $T = 5$ years, consisting of 255 weeks.

Other model parameters.

- Underlying resource model parameters: $\kappa = 0.50$, $\theta = \log(100)$, $\sigma_X = 0.50$.
- Discount rate: $\rho = 0.05$.

FIGURE 2 Model-implied distribution of the true reserve amount at $t = 0$ and $t = 2$ (variance reduction time frame), conditional on being at the initial estimate $\mu = v^{(L+1)}$, in the case of (a) slow learning and (b) fast learning.

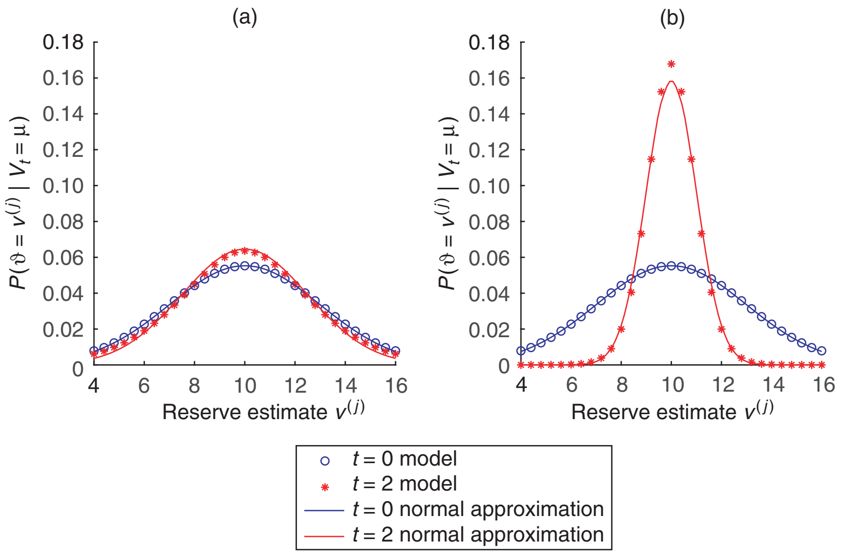

Lines correspond to normal approximations for the invariant distribution with parameters (μ, σ_0^2) and $(\mu, \sigma_{T'}^2)$, respectively.

- Extraction rate parameters: $\alpha = 1.00$, $\beta = 0.05$, $\gamma = 0.90$, $\varepsilon = 2.00$.
- Markov chain parameters: 31 states.

Further, we will consider the cases of slow and fast learning, which we model by changing the reduction in reserve estimate variance ($\sigma_{T'}^2 = 2.5 \times 10^8$ versus $\sigma_{T'}^2 = 10^8$) with $T' = 2$ years. We will also consider the case with running costs and the case without ($c = 20$ and $c = 0$, respectively). Finally, we compare these with the case where this is no learning by setting $a = 1$ and $b = 0$.

Figure 2 shows the effect of learning on the distribution of the true reserve amount. We find in both the slow- and fast-learning cases that the distribution (conditional on being at the initial reserve estimate) displays more of a peak around the mid-state. This reflects the fact that, by this time, the investor has learned more about the reserve amount. Given that they are in the mid-state at $t = 2$ (that is, the time frame given for calibrated variance reduction), they are more inclined to believe μ to be the true reserve amount. This effect is slight in the slow-learning case and more pronounced in the fast-learning one; this is because the investor can be more confident in the

FIGURE 3 Exercise boundary for different reserve estimate states when the agent's learning rate is (a) low and (b) high, and the agent incurs running costs.

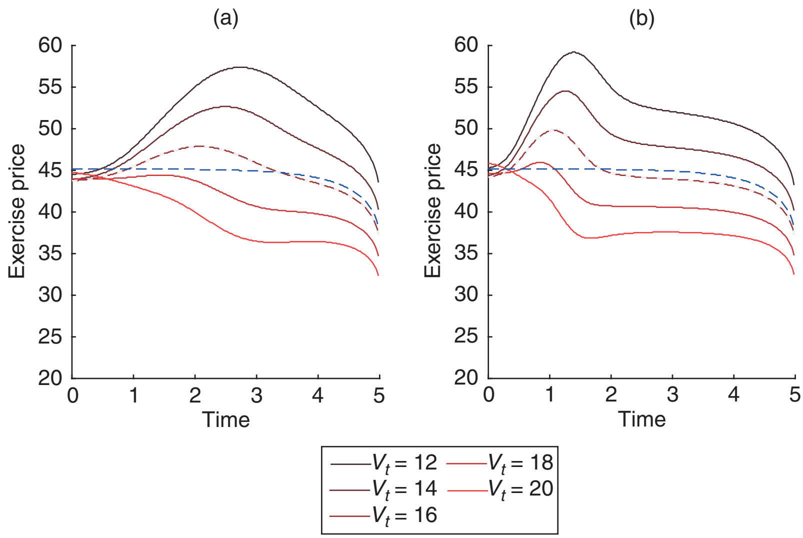

Lighter lines correspond to higher reserve estimates, the dashed red line corresponds to the mid-state (ie, the initial reserve estimate) and the blue line represents the exercise boundary when there is no learning.

latter case, having learned more about the reserve in the same amount of time. Note that the model-implied probabilities match the normal approximation only at $t = 0$ since this was the only constraint imposed on the calibration procedure.

Figure 3 shows the optimal exercise boundary for an agent whose learning happens at different rates for different volume estimates (ie, different states of the Markov chain). For simplicity, we display only a select number of states near the mid-state, as these represent the most relevant volume estimates for the investor. The y-axis of the figure shows the spot price of the commodity at which the investor would exercise their option, while the x-axis gives the time elapsed (measured in years). Figure 3 shows that as the agent's volume estimate increases, the exercise boundary shifts down; this can be explained by the fact that the larger the reserve, the lower the commodity spot price required to justify an investment.

We observe that the inclusion of learning, be it fast or slow, has a profound effect on the shape of the exercise boundary. In the no-learning case, this boundary is nonincreasing with time, whereas the slow- and fast-learning cases lead to exercise boundaries with nontrivial shapes, displaying increases and decreases at different rates as well as behavior that varies depending on the investor's prevailing reserve estimate.

FIGURE 4 Exercise boundary for different reserve estimate states when the agent's learning rate is (a) low and (b) high, and the agent does not incur any running costs.

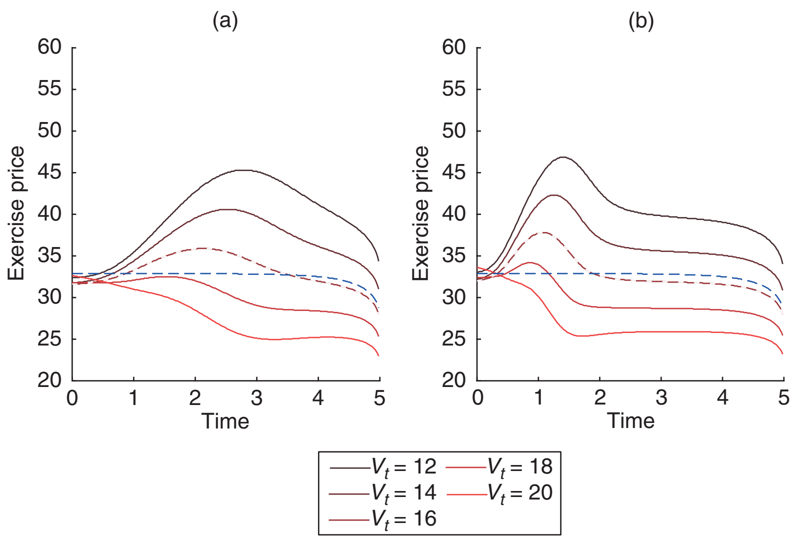

Lighter lines correspond to higher reserve estimates, the dashed red line corresponds to the mid-state (initial reserve estimate) and the blue line is the exercise boundary when there is no learning.

The explanation for the general shape taken by these curves is as follows. When the investor is in a low state (that is, their reserve estimate is low), the boundary is increased, allowing the learning process to potentially drive the estimate upward and leading to a higher project value in the event that the reserve estimate increases. Eventually, the investor decides that the learning process has provided enough information for them to feel confident that whatever state they currently occupy is in fact the true reserve amount. At this point, the exercise boundary goes back to its typical decreasing pattern. The time at which this "inflection point" occurs depends on the rate of learning; when learning is fast, the turnaround comes sooner.

Figure 4 shows the exercise boundaries for different volume states when the agent does not incur running costs. Observe that the exercise boundaries in both the fast- and slow-learning cases have the same general shape as the boundaries in Figure 3. Indeed, adding running costs gives the effect of an upward parallel shift of the exercise boundaries, as the commodity spot price must be higher to justify the investment after all running costs have been accounted for.

Figure 5 compares the value of the option at various points in time for the different learning rates, assuming that the investor's estimate is equal to the initial estimate.

FIGURE 5 Option value as a function of spot price through time assuming the reserve estimate is equal to the mid-state (initial estimate), $V_t = \mu$, with (a) slow learning and (b) fast learning.

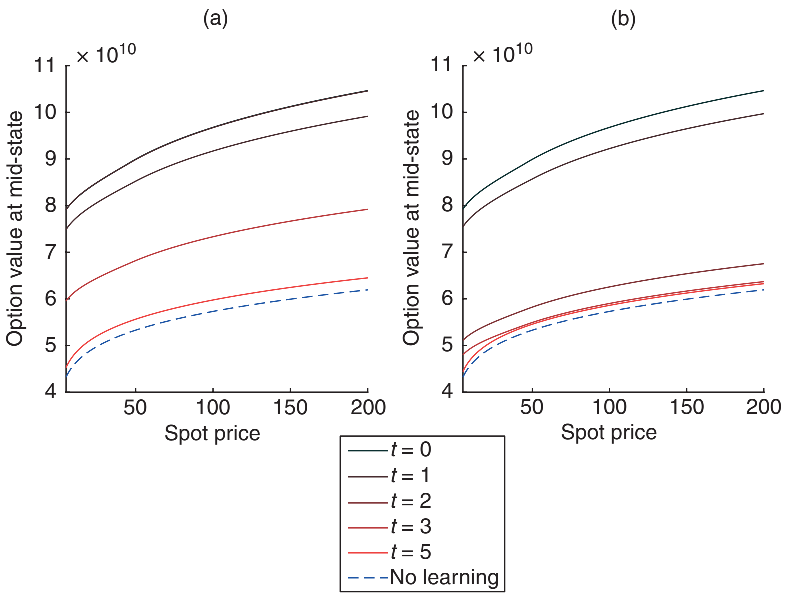

Dashed lines correspond to the option value at maturity when the investor does not learn.

We contrast this with the case where no learning occurs, which is very insensitive to time-to-maturity compared with the learning case.[5] We find that the value of the option to wait and learn is high at the beginning and gradually decreases as the expiry of the option approaches; this is because the agent is left with less and less time to learn. As expected, this decline in value is accelerated in the fast-learning case, as the investor becomes more confident about the reserve amount at an earlier point in time. Note that the effect of the passage of time on the option value is different in other states.

6 CONCLUSION

In this paper, we show how technical uncertainty can be incorporated into the decision to invest in a commodity reserve. This uncertainty stems from not knowing the

[5] Note that, since the value of the option changes very slightly with time-to-maturity in the no-learning case, we plot the no-learning option value curve at $t = 0$ only in Figure 5.

volume of the commodity stored in the reserve and is compounded by the uncertainty surrounding the value of the reserve (since future spot prices are unknown).

The agent has the option to "wait and see" before making an irreversible investment to exploit the commodity reserve. In our model, the agent learns about the volume of the commodity stored in the reserves as time goes by; thus, the option to delay investment is valuable because it allows the agent to learn more about the volume and to wait for the optimal market conditions (ie, the spot price of the commodity) before sinking their investment.

To model the reserve volume and the technical uncertainty, we adopt a continuous-time Markov chain. In our model, the agent learns about the volume in the reserve as time goes by; the accuracy of their estimates also improves with time. We show how to calculate the value of the option to delay investment and discuss the agent's optimal investment threshold.

We show how the exercise boundary depends on the agent's estimate of the volume (which is determined by the Markov chain state) along with the rate at which they refine their estimates of the reserve. For example, we show that when the option to invest is far from expiry and the volume estimate is low, the value attached to waiting and gathering more information is higher for an agent who can learn about the volume of the reserves quickly than for an agent who learns at a very slow rate.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

We are grateful to two anonymous referees for their valuable comments.

REFERENCES

Armstrong, M., Galli, A., Bailey, W., and Couët, B. (2004). Incorporating technical uncertainty in real option valuation of oil projects. *Journal of Petroleum Science and Engineering* **44**(1), 67–82 (https://doi.org/10.1016/j.petrol.2004.02.006).

Brennan, M. J., and Schwartz, E. S. (1985). Evaluating natural resource investments. *Journal of Business* **58**(2), 135–157 (https://doi.org/10.1086/296288).

Cartea, Á., and Figueroa, M. G. (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* **12**(4), 313–335 (https://doi.org/10.1080/13504860500117503).

Cartea, Á., and González-Pedraz, C. (2012). How much should we pay for interconnecting electricity markets? A real options approach. *Energy Economics* **34**(1), 14–30 (https://doi.org/10.1016/j.eneco.2011.06.002).

Cartea, Á., and Jaimungal, S. (2017). Irreversible investments and ambiguity aversion. *International Journal of Theoretical and Applied Finance* **20**(07), 1750044 (https://doi.org/10.1142/S0219024917500443).

Cartea, Á., and Williams, T. (2008). UK gas markets: the market price of risk and applications to multiple interruptible supply contracts. *Energy Economics* **30**(3), 829–846 (https://doi.org/10.1016/j.eneco.2007.03.001).

Cartea, Á., Jaimungal, S., and Zhen, Q. (2016). Model uncertainty in commodity markets. *SIAM Journal on Financial Mathematics* **7**(1), 1–33 (https://doi.org/10.1137/15M1027243).

Cartea, Á., Donnelly, R. F., and Jaimungal, S. (2017). Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics* **8**(1), 635–671 (https://doi.org/10.1137/16M106282X).

Cartea, Á., Donnelly, R. F., and Jaimungal, S. (2018). Portfolio liquidation and ambiguity aversion. In *High-Performance Computing in Finance: Problems, Methods, and Solutions*, Dempster, M. A. H., Kanniainen, J., Keane, J., and Vynckier, E. (eds). Chapman and Hall, London.

Cortázar, G., Schwartz, E. S., and Casassus, J. (2001). Optimal exploration investments under price and geological-technical uncertainty: a real options model. *R&D Management* **31**(2), 181–189.

Coulon, M., Powell, W. B., and Sircar, R. (2013). A model for hedging load and price risk in the Texas electricity market. *Energy Economics* **40**, 976–988 (https://doi.org/10.1016/j.eneco.2013.05.020).

Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of Political Economy* **97**(3), 620–638 (https://doi.org/10.1086/261619).

Dixit, A., and Pindyck, R. (1994). *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

Elliott, R. J., Miao, H., and Yu, J. (2009). Investment timing under regime switching. *International Journal of Theoretical and Applied Finance* **12**(04), 443–463 (https://doi.org/10.1142/S0219024909005361).

Fleten, S.-E., Heggedal, A. M., and Siddiqui, A. (2011). Transmission capacity between Norway and Germany: a real options analysis. *The Journal of Energy Markets* **4**(1), 121–147 (https://doi.org/10.21314/JEM.2011.054).

Himpler, S., and Madlener, R. (2014). Optimal timing of wind farm repowering: a two-factor real options analysis. *The Journal of Energy Markets* **7**(3), 3–34 (https://doi.org/10.21314/JEM.2014.111).

Jackson, K., Jaimungal, S., and Surkov, V. (2008). Fourier space time stepping for option pricing with Lévy models. *The Journal of Computational Finance* **12**(2), 1–29 (https://doi.org/10.21314/JCF.2008.178).

Jaimungal, S., and Lawryshyn, Y. (2015). Incorporating managerial information into real option valuation. In *Commodities, Energy and Environmental Finance*, pp. 213–238. Springer (https://doi.org/10.1007/978-1-4939-2733-3_8).

Jaimungal, S., and Surkov, V. (2011). Lévy-based cross-commodity models and derivative valuation. *SIAM Journal on Financial Mathematics* **2**(1), 464–487 (https://doi.org/10.1137/100791609).

Jaimungal, S., and Surkov, V. (2013). Valuing early-exercise interest-rate options with multi-factor affine models. *International Journal of Theoretical and Applied Finance* **16**(06), 1350034 (https://doi.org/10.1142/S0219024913500349).

Jaimungal, S., De Souza, M. O., and Zubelli, J. P. (2013). Real option pricing with mean-reverting investment and project value. *European Journal of Finance* **19**(7–8), 625–644 (https://doi.org/10.1080/1351847X.2011.601660).

Kiesel, R., Schindlmayr, G., and Boerger, R. H. (2009). A two-factor model for the electricity forward market. *Quantitative Finance* **9**(3), 279–287 (https://doi.org/10.1080/14697680802126530).

Kobari, L., Jaimungal, S., and Lawryshyn, Y. (2014). A real options model to evaluate the effect of environmental policies on the oil sands rate of expansion. *Energy Economics* **45**, 155–165 (https://doi.org/10.1016/j.eneco.2014.06.010).

McDonald, R., and Siegel, D. (1985). Investment and the valuation of firms when there is an option to shut down. *International Economic Review* **26**(2), 331–349 (https://doi.org/10.2307/2526587).

McDonald, R., and Siegel, D. (1986). The value of waiting to invest. *Quarterly Journal of Economics* **101**(4), 707–728 (https://doi.org/10.2307/1884175).

Metcalf, G., and Hasset, K. (1995). Investment under alternative return assumptions: comparing random walk and mean reversion. *Journal of Economic Dynamics and Control* **19**(8), 1471–1488 (https://doi.org/10.1016/0165-1889(94)00838-9).

Pindyck, R. S. (1980). Uncertainty and exhaustible resource markets. *Journal of Political Economy* **88**(6), 1203–1225 (https://doi.org/10.1086/260935).

Sadowsky, J. R. (2005). The value of learning in the product development stage: a real options approach. Working Paper 721597, Social Science Research Network.

Sarkar, S. (2003). The effect of mean reversion on investment under uncertainty. *Journal of Economic Dynamics and Control* **28**, 377–396 (https://doi.org/10.1016/S0165-1889(02)00181-1).

Taschini, L., and Urech, S. (2010). The real option to fuel switch in the presence of expected windfall profits under the EU ETS. *The Journal of Energy Markets* **3**(2), 27–47 (https://doi.org/10.21314/JEM.2010.042).

Trigeorgis, L. (1996). *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. MIT Press, Cambridge, MA.

Tsekrekos, A. E. (2010). The effect of mean reversion on entry and exit decisions under uncertainty. *Journal of Economic Dynamics and Control* **34**(4), 725–742 (https://doi.org/10.1016/j.jedc.2009.10.015).

Weron, R. (2007). *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach.* Wiley Finance Series, Volume 403. Wiley.