SIAM J. SCI. COMPUT.
Vol. 30, No. 1, pp. 339–368

A SEMI-LAGRANGIAN APPROACH FOR NATURAL GAS STORAGE VALUATION AND OPTIMAL OPERATION*

ZHULIANG CHEN† AND PETER A. FORSYTH†

Abstract. The valuation of a gas storage facility is characterized as a stochastic control problem, resulting in a Hamilton–Jacobi–Bellman (HJB) equation. In this paper, we present a semi-Lagrangian method for solving the HJB equation for a typical gas storage valuation problem. The method is able to handle a wide class of spot price models that exhibit mean-reverting seasonality dynamics and price jumps. We develop fully implicit and Crank–Nicolson timestepping schemes based on a semi-Lagrangian approach and prove the convergence of fully implicit timestepping to the viscosity solution of a modified HJB equation posed on a bounded domain, provided that a strong comparison result holds. The semi-Lagrangian approach avoids Policy-type iterations required by an implicit finite difference method without requiring additional cost. Numerical experiments are presented for several variants of the basic scheme.

Key words. gas storage, semi-Lagrangian, HJB equation, stochastic control, viscosity solution

AMS subject classifications. 65N06, 93C20

DOI. 10.1137/060672911

1. Introduction.

Similar to other commodities such as fuel and electricity, natural gas prices exhibit seasonality dynamics due to fluctuations in demand [27]. Natural gas storage facilities are constructed to provide a cushion for such fluctuations by releasing natural gas from storage in seasons with high demand.

Recently, several authors [1, 31, 34, 35, 25, 11, 8] have discussed the no-arbitrage value of natural gas storage facilities (or, equivalently, the values of gas storage contracts for leasing the storage facilities). The value of a gas storage facility can be regarded as the maximum expected revenues under the risk-neutral measure that the operator of the facility can obtain by optimally operating the facility, that is, "buying low" and "selling high." As a result, the valuation of gas storage facilities is characterized as a stochastic control problem.

Two contrasting approaches have been used to solve this problem numerically: partial differential equation (PDE)-based approaches and simulation-based methods. Assuming that the control is of bang-bang type (that is, the control takes values only from a finite set), simulation-based methods [11, 8] can be used to directly solve the stochastic control problem. These methods are well suited for solving multi-dimensional problems (e.g., problems with four state variables), which the PDE-based approaches cannot handle. Nevertheless, it is known that such methods have difficulty achieving high accuracy. Furthermore, if the control is not of bang-bang type, simulation-based methods will have to approximate the control as piecewise constant, which will substantially increase the computational cost. See [32] for descriptions of control problems which are not of the bang-bang type for valuation of electricity power plants.

*Received by the editors October 20, 2006; accepted for publication (in revised form) June 4, 2007; published electronically December 21, 2007. This work was supported by the Natural Sciences and Engineering Research Council of Canada, and by a Morgan Stanley Equity Market Microstructure Research Grant. The views expressed herein are solely those of the authors, and not those of any other person or entity, including Morgan Stanley.

http://www.siam.org/journals/sisc/30-1/67291.html

†David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada N2L 3G1 (z4chen@uwaterloo.ca, paforsyt@uwaterloo.ca).

As shown in [26, 22], the value of a stochastic control problem is identical to the viscosity solution of a Hamilton–Jacobi–Bellman (HJB) equation. As a result, the value of a gas storage contract can be computed by solving the corresponding HJB equation using PDE-based approaches. In general, the solution to the HJB equation may not be unique. As noted in [23, 2], it is important to ensure that a numerical scheme converges to the viscosity solution of the equation, which is the appropriate value of the corresponding stochastic control problem.

In [31], an explicit finite difference scheme is used to price gas storage contracts. As is well known, explicit timestepping suffers from timestep restrictions due to stability considerations. Alternatively, the authors of [23] present implicit finite difference schemes, which eliminate the timestep restriction, for solving general controlled HJB equations. However, these schemes require solution of nonlinear discretized algebraic equations using a Policy iteration at each timestep. Reference [23] introduces another type of implicit scheme that approximates the control as piecewise constant to avoid the need for solving nonlinear equations, at the expense of solving a number of linear problems at each timestep. Similar to the simulation-based methods, if the control is not of bang-bang type, such schemes will be computationally expensive. In [34], a finite element semi-Lagrangian scheme is developed to solve an HJB equation with no bang-bang control for certain gas storage contracts. In [35, 25], a wavelet method coupled with a semi-Lagrangian approach is used to solve the gas storage HJB equation. While the wavelet method shows promise, it is difficult to obtain a rigorous proof of convergence to the viscosity solution.

In this paper, following [31], we formulate the value of a gas storage facility as a two-dimensional HJB equation. We introduce a modified HJB equation on a bounded computational domain. This modified equation has the advantage that no additional boundary conditions are required on the artificial boundaries, which is consistent with the original problem posed on an unbounded domain. For the particular model selected, the controls can be shown to be of the bang-bang type [31]. Since the equation is convection dominated, that is, the equation has no diffusion in one of the coordinate directions, it is well known to be difficult to solve numerically. Initially introduced by [19, 28] for atmospheric and weather numerical predictions, semi-Lagrangian schemes can effectually reduce the numerical problems arising for convection dominated equations. In [17], a semi-Lagrangian method is applied to price continuously observed American Asian options.

We apply the semi-Lagrangian approach to solve the HJB equation for gas storage contracts in this paper. We consider only a fairly simple form for the spot price model, but our method can easily generalize to more complex spot price models. To illustrate this, we also include numerical experiments using a process that incorporates price jumps.

The semi-Lagrangian discretizations require solution of discrete local optimization problems in order to determine optimal control values. Although the controls for the exact solution of the pricing HJB equation are of bang-bang type, the solutions of the discrete optimization problems may allow controls which are not possible controls for the exact solution. Consequently, two methods are used to solve the optimization problems: the *bang-bang* method that considers only bang-bang controls as appearing in the exact solution, and the *no bang-bang* method that actually solves the discrete problem and hence allows any admissible controls.

The main results of this paper are summarized as follows:

- We demonstrate that a semi-Lagrangian method can be used to price gas storage contracts and obtain optimal control strategies under a wide range

of spot price models that exhibit mean-reverting seasonality dynamics and price jumps. A semi-Lagrangian method reduces the problem to solving a decoupled set of one-dimensional nonlinear discrete PDEs at each timestep. This makes implementation of this method very straightforward in a software library which is capable of pricing path-dependent options [36].

- We show that, as opposed to implicit finite difference discretizations, the semi-Lagrangian discretizations avoid the need for Policy-type iterations at each timestep. Effectively, we replace the Policy iteration by an operation involving a local optimization over known solution values from the previous timestep. The solution of this local optimization problem can be obtained cheaply, at the expense of an error proportional to the mesh size.
- Provided that a strong comparison property holds for the PDE, we prove that the fully implicit, semi-Lagrangian discretization converges to the viscosity solution of the modified HJB equation on a bounded domain by verifying the stability, consistency, and monotonicity of the scheme and using the basic results in [7, 2]. We pay particular attention to the boundary conditions and properties of the admissible controls on the boundary of the computational domain, to ensure that the conditions required for convergence to the viscosity solution are met.
- Numerical results show that fully implicit and Crank–Nicolson timestepping schemes using the semi-Lagrangian method converge to the same solution. In the case of fully implicit timestepping, the no bang-bang method converges smoothly at a first-order rate for some cases that cause numerical difficulties for the bang-bang method. The no bang-bang method is only about 10 percent more costly, in terms of CPU time, compared to the bang-bang method when tested using a typical example. Although most gas storage models in the literature have controls of the bang-bang type, the no bang-bang method is more general and can be applied to a wide class of optimal control problems where the converged controls may not be of bang-bang type.
- Numerical results also indicate that Crank–Nicolson timestepping does not appear to converge at a higher than first-order rate, and hence fully implicit timestepping is probably a better choice since it guarantees convergence to the viscosity solution, provided that a strong comparison result holds. Fully implicit timestepping is also straightforward to implement.

2. Gas storage equation.

This section defines the natural gas storage problem and corresponding PDE for the value of a gas storage contract. The section is arranged as follows: first, we define some notation for the problem; we then present a mean-reverting model for the natural gas spot price. We give the PDE for the value of a natural gas storage contract with respect to the gas spot price model. Finally, we completely specify the gas storage problem by providing the boundary conditions for the PDE.

2.1. Problem notation.

We use the following notation for the natural gas storage problem:

- P: the current spot price per unit of natural gas.
- I: current amount of working gas inventory. We assume that I can be any value lying within the domain $[0, I_{\max}]$.
- $\hat{V}(P, I, t)$: value of the natural gas storage facility (e.g., the leasing rate of the facility) with respect to natural gas price P and inventory level I at time t.

- T: expiry time of the contract.
- c: control variable that represents the rate of producing gas from or injecting gas into the gas storage ($c > 0$ represents production, $c < 0$ represents injection). If $c = 0$, then no operation is performed on the storage.
- $c_{\max}(I)$: the maximum rate at which gas can be released from storage as a function of inventory levels, $c_{\max}(I) > 0$. We use the expression in [31],

$$c_{\max}(I) = k_1\sqrt{I}, \tag{2.1}$$

  where k_1 is a positive constant. This implies $c_{\max}(0) = 0$ with the physical meaning that no gas can be produced if the gas storage is empty.
- $c_{\min}(I)$: $|c_{\min}(I)|$ is the maximum rate at which gas can be injected into storage as a function of inventory levels. Note that $c_{\min}(I) < 0$, with our sign convention that $c > 0$ represents production. We use the expression from [31],

$$c_{\min}(I) = -k_2\sqrt{\frac{1}{I + k_3} - \frac{1}{k_4}}, \tag{2.2}$$

  where k_2, k_3, and k_4 are positive constants and k_2, k_3, k_4 satisfy the constraint $c_{\min}(I_{\max}) = 0$, which means that no gas injection is possible if the gas storage is full. Equation (2.2) implies that

$$|c_{\min}(I)| \le \left|\text{const.}\sqrt{I_{\max} - I}\right|, \qquad I < I_{\max}, \quad I \to I_{\max}. \tag{2.3}$$

- $a(I, c)$: the rate of gas loss given a gas production rate of c, when I units are currently in storage. In general, the change in gas inventory satisfies

$$\frac{dI}{dt} = -(c + a(I, c)), \tag{2.4}$$

  where usually $a(I, c) \ge 0$. We use the model in [31],

$$a(I, c) = a(c) = \begin{cases} 0 & \text{if } c \ge 0 \text{ (producing gas)}, \\ k_5 & \text{if } c < 0 \text{ (injecting gas)}, \end{cases} \tag{2.5}$$

  where $k_5 > 0$ is a constant.

Note that if we are using a control c satisfying $-k_5 < c < 0$, that is, injecting at a rate smaller than the rate of gas loss, then (2.5) implies that $c + a(c) > 0$. According to (2.4), this means that injecting natural gas into the gas storage decreases the gas inventory, which is unreasonable. Consequently, we further require that control c satisfy the constraint $c \in [c_{\min}(I), -k_5] \cup [0, c_{\max}(I)]$ so that

$$c + a(c) \le 0 \quad \text{if } c < 0 \text{ (injecting gas)}. \tag{2.6}$$

In other words, the operator of the gas storage facility is not allowed to inject and at the same time decrease the gas inventory. We point out that the constraint on the control also makes the boundary conditions well defined (this will be discussed in more detail in a subsequent section). For future reference, given any $I \in [0, I_{\max}]$, we define the set $C(I)$ as

$$C(I) = [c_{\min}(I), -k_5] \cup [0, c_{\max}(I)], \tag{2.7}$$

where we adopt the convention that $[\alpha, \beta] = \emptyset$ if $\alpha > \beta$.

2.2. Natural gas spot price model. In this subsection, we present a one-factor mean-reverting process for natural gas spot price. In the following we will refer to this process as a pure mean-reverting process. This process is able to capture the mean-reverting and seasonality effects of natural gas spot prices. However, there is a certain amount of controversy surrounding the precise form of a reasonable natural gas spot price model [27]. The numerical methods that we derive in this paper easily generalize to more complex spot price models.

Pure mean-reverting process. We note that the gas storage problem is a real options valuation problem [18, 33]. Hence, the usual derivation of the pricing equation starts with the actual spot price process, and a hedging argument generates a PDE with a risk-adjusted drift term [33]. We refer interested readers to [33, 27] for details. We will take a short cut here and simply consider the risk-adjusted price process, since this is the process which appears in the final PDE.

In this model, the risk-adjusted gas spot price follows a mean-reverting stochastic process with the equilibrium price varying over time to produce a seasonality effect. The risk-adjusted (or risk-neutral) gas spot price is modeled by a stochastic differential equation given by

$$dP = \alpha(K(t) - P)dt + \hat{\sigma}(P)PdZ, \tag{2.8}$$

$$K(t) = K_0 + \beta_{SA}\sin(4\pi(t - t_{SA})), \tag{2.9}$$

where

- $\alpha > 0$ is the mean-reverting rate,
- $K(t) \geq 0$ is the long-term equilibrium price that incorporates seasonality,
- $\hat{\sigma}(P)$ is the volatility as a function of P,
- dZ is an increment of the standard Gauss–Wiener process,
- $K_0 \geq 0$ is the equilibrium price without seasonality effect,
- β_{SA} is the semiannual seasonality parameter,
- t_{SA} is the seasonality centering parameters, representing the time of semiannual peak of equilibrium price in summer and winter.

According to (2.9), the equilibrium price $K(t)$ is a periodic function with period $1/2$. This models the seasonal evolution of the annual equilibrium price; e.g., $K(t)$ exhibits two peaks annually, respectively corresponding to high natural gas spot prices in summer and winter.

2.3. The PDE for natural gas storage facilities. For a financial contract such as the natural gas storage contract, a terminal payoff is given at the maturity $t = T$. In order to compute the value of the contract today, we need to solve the pricing PDE backwards in time from $t = T$ to $t = 0$. Let $\tau = T - t$ denote the current time-to-maturity. For ease of exposition, we will write our PDE in terms of τ so that we will solve the pricing PDE from $\tau = 0$ to $\tau = T$. Let $V(P, I, \tau)$ denote the value of a natural gas storage facility as a function of (P, I, τ). In terms of the facility value $\hat{V}(P, I, t)$ at forward times with respect to t, we have the identity $V(P, I, \tau) = \hat{V}(P, I, T - \tau) = \hat{V}(P, I, t)$.

Based on standard hedging arguments in the financial valuation literature [31, 34, 35], the value of a natural gas storage facility in terms of τ, assuming that the risk-adjusted natural gas spot price follows the stochastic process defined in (2.8), is

given by

(2.10)
$$V_\tau = \frac{1}{2}(\hat{\sigma}(P)P)^2 V_{PP} + \alpha(K(t) - P)V_P + \max_{c \in C(I)} \{(c - a(c))P - (c + a(c))V_I\} - rV,$$

where $C(I)$ is given in (2.7) and $r > 0$ is the continuously compounded risk-free interest rate. The nonlinear PDE (2.10) is known as the HJB equation. In (2.10), the term $(c - a(c))P$ represents the instantaneous rate of revenue obtained by producing natural gas from $(c > 0)$ or injecting gas into $(c < 0)$ the gas storage facility. The term $-(c + a(c))V_I$ results from the rate of inventory change (recall (2.4)).

2.4. Boundary conditions for the pricing PDE. In order to completely specify the gas storage problem, we need to provide boundary conditions. We will proceed in an informal fashion here to derive plausible boundary conditions. These boundary conditions will be justified in a more rigorous fashion in section 4.

A number of terminal boundary conditions can be used. Note that since we will be solving backwards in time, the terminal state occurs at $\tau = 0$. Typical examples include the following:

- a zero payoff, as suggested by [31]: $V(P, I, \tau = 0) = 0$.
- a nonnegative payoff obtained by selling all the leftover gas in the storage at the maximum rate; that is, $V(P, I, \tau = 0)$ is the solution $\bar{V}$ to the PDE obtained by fixing control $c = c_{\max}(I)$ in PDE (2.10) and solving the resulting PDE backwards from $\tau = 0$ to $\tau = \infty$ with $\bar{V}(P, I, \tau = 0) = 0$. We then specify $V(P, I, \tau = 0) = \bar{V}(P, I, \tau = \infty)$.
- the penalty payoff introduced by [11]:

  (2.11)
  $$V(P, I, \tau = 0) = \text{const.} \cdot P \cdot \min\left(I(t = T) - I(t = 0), 0\right),$$

  where const. > 0. This has the meaning that a penalty will be charged if the gas inventory in storage when the facility is returned is less than the gas inventory at the inception of the contract.

The domain for the PDE (2.10) is $P \times I \in [0, \infty] \times [0, I_{\max}]$. For computational purposes, we need to solve the PDE in a finite computational domain $[0, P_{\max}] \times [0, I_{\max}]$.

As $I \to 0$, from (2.1)–(2.7) we have that

(2.12)
$$c + a(c) \le 0 \quad \forall c \in C(I),\ I \to 0.$$

Hence the characteristics are outgoing in the I direction at $I = 0$, and we simply solve the PDE along the $I = 0$ boundary; no further information is needed. Condition (2.12) has the interpretation that gas cannot be produced from a facility which is empty.

Similarly, as $I \to I_{\max}$, equations (2.1)–(2.7) imply that

(2.13)
$$c + a(c) \ge 0 \quad \forall c \in C(I),\ I \to I_{\max},$$

which again means that the characteristics are outgoing in the I direction at $I = I_{\max}$. Consequently, we simply solve the PDE along the $I = I_{\max}$ boundary; no further information is needed. Condition (2.13) has the interpretation that gas cannot be injected into the storage facility when it reaches full capacity.

Taking the limit of (2.10) as $P \to 0$, we obtain

(2.14)
$$V_\tau = \alpha K(t) V_P + \max_{c \in C(I)} \{-(c + a(c))V_I\} - rV.$$

Since $\alpha K(t) \geq 0$, we can solve (2.14) without requiring additional boundary conditions, as we do not need information from outside the computational domain $[0, P_{\max}]$.

In this paper, we assume that $\hat{\sigma}(P)$ is a continuous function that satisfies $\hat{\sigma}(P) = \sigma$ for $P \in [0, P_{\max} - \epsilon]$, where $\sigma > 0$ is a constant and $\epsilon > 0$ is a constant close to zero. For $P \in [P_{\max} - \epsilon, P_{\max}]$, $\hat{\sigma}(P)$ approaches 0 continuously as $P \to P_{\max}$. In other words, the volatility is constant for most values of P and decreases to zero as $P \to P_{\max}$. The decreasing behavior of $\hat{\sigma}(P)$ at the far boundary $P = P_{\max}$ generates a negligible error by choosing $P_{\max}$ sufficiently large (see section 5 for numerical results on different choices of $P_{\max}$).

Based on this form of $\hat{\sigma}(P)$, taking the limit of (2.10) as $P \to P_{\max}$, we have

$$V_\tau = \alpha(K(t) - P)V_P + \max_{c \in C(I)} \{(c - a(c))P - (c + a(c))V_I\} - rV, \quad P \to P_{\max}. \tag{2.15}$$

We will always choose $P_{\max} \gg K(t)$, and hence (2.15) can be solved at $P = P_{\max}$ without additional information.

The purpose of introducing the continuous function $\hat{\sigma}(P)$ is so that the boundary equation is the limit of the PDE from the domain interior. This will reduce the technical manipulations required to prove convergence of our numerical scheme to the viscosity solution of the modified gas storage problem (2.10)–(2.15) (see section 4). Since $\epsilon \ll 1$, the numerical implementation assuming that $\hat{\sigma}(P)$ has the above behavior is, for all practical purposes, the same as an implementation assuming that $\hat{\sigma}(P) = \sigma$ for $P < P_{\max}$ and $\hat{\sigma}(P) = 0$ for $P = P_{\max}$.

3. Semi-Lagrangian discretization. A semi-Lagrangian approach is presented in [17] for pricing continuously observed American Asian options under jump diffusion. In this section, we extend the semi-Lagrangian method in [17] to solve the HJB equation (2.10) that involves optimal control. The main idea used to construct a semi-Lagrangian discretization of the PDE (2.10) is to integrate the PDE along a semi-Lagrangian trajectory (defined below). Various semi-Lagrangian discretizations can be obtained by evaluating the resulting integrals using different numerical integration methods: for example, using the rectangular rule provides a fully implicit timestepping scheme, while using the trapezoidal rule gives a Crank–Nicolson timestepping scheme.

This section is arranged as follows: we first present an intuitive idea for developing a semi-Lagrangian discretization for (2.10). We then present both a fully implicit and a Crank–Nicolson timestepping scheme based on this idea. At the end of this section, we reformulate the discrete equations resulting from the semi-Lagrangian discretization into a matrix form as an aid in the analysis presented in later sections.

Prior to presenting the timestepping schemes, we introduce the following notation. We use an unequally spaced grid in P coordinates for the PDE discretization, represented by $[P_1, P_2, \ldots, P_{i_{\max}}]$. Similarly, we use an unequally spaced grid in the I direction denoted by $[I_1, I_2, \ldots, I_{j_{\max}}]$. We denote by $0 = \tau^0 < \cdots < \tau^N = T$ the discrete timesteps used to discretize the PDE (2.10). Let $\Delta\tau^n = \tau^{n+1} - \tau^n$. Let $V(P_i, I_j, \tau^n)$ denote the exact solution of (2.10) when the natural gas spot price resides at node P_i, the gas inventory at node I_j, and discrete time at τ^n. Let $V_{i,j}^n$ denote an approximation of the exact solution $V(P_i, I_j, \tau^n)$. Let $\mathcal{L}$ be a differential operator represented by

$$\mathcal{L}V = \frac{1}{2}(\hat{\sigma}(P)P)^2 V_{PP} + \alpha(K(t) - P) - rV. \tag{3.1}$$

Using the differential operator (3.1), we can rewrite natural gas storage pricing PDE (2.10) as

$$\min_{c\in C(I)} \{V_\tau + (c + a(c))V_I - (c - a(c))P - \mathcal{L}V\} = 0. \tag{3.2}$$

We use standard finite difference methods to discretize the operator $\mathcal{L}V$. Let $(\mathcal{L}_h V)^n_{i,j}$ denote the discrete value of the differential operator (3.1) at node (P_i, I_j, τ^n). The operator (3.1) can be discretized using central, forward, or backward differencing in the P, I directions to give

$$(\mathcal{L}_h V)^n_{i,j} = \alpha^n_i V^n_{i-1,j} + \beta^n_i V^n_{i+1,j} - (\alpha^n_i + \beta^n_i + r)V^n_{i,j}, \tag{3.3}$$

where α^n_i and β^n_i are determined using the algorithm in Appendix A. The algorithm guarantees that α^n_i and β^n_i satisfy the following positive coefficient condition:

$$\alpha^n_i \geq 0, \quad \beta^n_i \geq 0, \qquad i = 1, \ldots, i_{\max},\ j = 1, \ldots, j_{\max},\ n = 1, \ldots, N. \tag{3.4}$$

As we will demonstrate in section 4, the positive coefficient property (3.4) is a sufficient condition to ensure convergence of a semi-Lagrangian fully implicit timestepping scheme to the viscosity solution of the HJB equation (2.10).

Now we give the intuition for developing the semi-Lagrangian discretization schemes. Let us consider a path (or a semi-Lagrangian trajectory) $\mathcal{I}(\tau)$ that follows the ODE

$$\frac{d\mathcal{I}}{d\tau} = c + a(c). \tag{3.5}$$

According to (3.5), we can write the term $V_\tau + (c + a(c))V_I$ in (3.2) in the form of a Lagrangian directional derivative

$$\frac{DV}{D\tau} = \frac{\partial V}{\partial \tau} + \frac{\partial V}{\partial I}\frac{d\mathcal{I}}{d\tau}. \tag{3.6}$$

Then (3.2) can be rewritten as

$$\min_{c\in C(I)} \left\{ \frac{DV}{D\tau} - (c - a(c))P - \mathcal{L}V \right\} = 0. \tag{3.7}$$

Let us analyze (3.5)–(3.7) from a discrete point of view, that is, consider discrete grid points and discrete times. Let $\mathcal{I}\big(\tau; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau)\big)$ denote a path satisfying (3.5), which arrives at a discrete grid point (P_i, I_j) at $\tau = \tau^{n+1}$ for P_i being held constant and control following a path $\zeta_{i,j}(\tau)$ with respect to τ. Let $\mathcal{I}\big(\tau^n; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau^n)\big)$ be the departure point of this path at $\tau = \tau^n$, which can be computed by solving

$$\begin{cases} \dfrac{d\mathcal{I}}{d\tau}\big(\tau; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau)\big) = \zeta_{i,j}(\tau) + a\big(\zeta_{i,j}(\tau)\big) & \text{for } \tau < \tau^{n+1}, \\ \mathcal{I}\big(\tau; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau)\big) = I_j & \text{for } \tau = \tau^{n+1}, \end{cases} \tag{3.8}$$

from $\tau = \tau^{n+1}$ to $\tau = \tau^n$. We can write the solution of (3.8) in the integral form as

$$\mathcal{I}\big(\tau = \tau^n; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau = \tau^n)\big) = I_j - \int_{\tau^n}^{\tau^{n+1}} \Big[\zeta_{i,j}(\tau) + a\big(\zeta_{i,j}(\tau)\big)\Big] d\tau. \tag{3.9}$$

Note that from (3.9), the departure point $\mathcal{I}\big(\tau^n; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau^n)\big)$ will not necessarily coincide with a grid point in the I direction. To simplify the notation, in the following we will use $\mathcal{I}(\tau) = \mathcal{I}\big(\tau; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau)\big)$ without causing ambiguity.

Integrating both sides of (3.7) along the path $\mathcal{I}(\tau)$ from $\tau = \tau^n$ to $\tau = \tau^{n+1}$, with P fixed at P_i and control variable c following the path $\zeta_{i,j}(\tau)$, gives

$$\int_{\tau^n}^{\tau^{n+1}} \min_{\zeta_{i,j}(\tau)\in C(\mathcal{I}(\tau))} \left\{ \frac{DV}{D\tau}\big(P_i, \mathcal{I}(\tau), \tau\big) - \big(\zeta_{i,j}(\tau) - a(\zeta_{i,j}(\tau))\big) P_i - \mathcal{L}V\big(P_i, \mathcal{I}(\tau), \tau\big) \right\} d\tau = 0. \tag{3.10}$$

Interchanging the integral and the min operator in (3.10), assuming that they are interchangeable, and using the identity

$$\int_{\tau^n}^{\tau^{n+1}} \frac{DV}{D\tau}\big(P_i, \mathcal{I}(\tau), \tau\big) d\tau = V\big(P_i, I_j, \tau^{n+1}\big) - V\big(P_i, \mathcal{I}(\tau^n), \tau^n\big), \tag{3.11}$$

we obtain

$$V\big(P_i, I_j, \tau^{n+1}\big) = \min_{\zeta_{i,j}(\tau)\in C(\mathcal{I}(\tau))} \left\{ V\left(P_i, \mathcal{I}(\tau^n), \tau^n\right) + \int_{\tau^n}^{\tau^{n+1}} \big(\zeta_{i,j}(\tau) - a\big(\zeta_{i,j}(\tau)\big)\big) P_i d\tau + \int_{\tau^n}^{\tau^{n+1}} \mathcal{L}V\big(P_i, \mathcal{I}(\tau), \tau\big) d\tau \right\}, \tag{3.12}$$

where $\mathcal{I}(\tau) = \mathcal{I}\big(\tau; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau)\big)$.

Remark 3.1 (interchanging the order of operations in (3.10)). The integral and the min operators may not be interchangeable. Moreover, the derivatives in (3.10) may not exist since the value function V may not be smooth and needs to be considered in the sense of the viscosity solution. Thus, our derivation is not rigorous. However, our purpose here is to illustrate the main idea for developing the schemes. The rigorous proof of the convergence of the semi-Lagrangian fully implicit discretization to the viscosity solution of (2.10) will be given in section 4.

By evaluating the integrals in (3.9), the departure point of the semi-Lagrangian trajectory, and value function (3.12) using various numerical integration schemes, we are able to obtain semi-Lagrangian discretizations of different orders in time. In this paper, we will present the fully implicit and Crank–Nicolson timestepping schemes, which result from approximating the integrals using the rectangular rule and trapezoidal rule, respectively. We will use an approach similar to that suggested in [20].

3.1. Fully implicit timestepping. In the case of fully implicit timestepping, we approximate the integral in (3.9) using the rectangular rule at $\tau = \tau^{n+1}$. In other words, we evaluate (3.9) by assuming that

$$\zeta_{i,j}(\tau) \simeq \zeta_{i,j}(\tau^{n+1}) \quad \text{for } \tau \in [\tau^n, \tau^{n+1}]. \tag{3.13}$$

It is perhaps not immediately obvious why we evaluate $\zeta_{i,j}(\tau)$ at τ^{n+1} in approximation (3.13). In [13, Appendix B] we show that this choice corresponds to a discretization based on assuming that the operator of the facility can adjust the controls only at finite intervals. This assumes that the operator of the facility maximizes his risk-neutral expected value over the infinitesimal period of the cash flows. Note that it is also possible to derive PDE (2.10) using a hedging argument [34, 35]. In this case, we can also interpret the discretization as an application of no-arbitrage jump conditions at each discrete control observation time.

Let $\zeta_{i,j}^{n+1} = \zeta_{i,j}(\tau^{n+1})$ and let $\mathcal{I}_{j(i,n+1)}^n$ denote an approximation to $\mathcal{I}(\tau^n) = \mathcal{I}\big(\tau^n; P_i, I_j, \tau^{n+1}, \zeta_{i,j}(\tau^n)\big)$, the departure point of the semi-Lagrangian trajectory (3.8). The rectangular approximation of (3.9), assuming (3.13), gives

$$\mathcal{I}_{j(i,n+1)}^n = I_j - \Delta\tau^n\big(\zeta_{i,j}^{n+1} + a\big(\zeta_{i,j}^{n+1}\big)\big), \tag{3.14}$$

where $\Delta\tau^n = \tau^{n+1} - \tau^n$.

The control $\zeta_{i,j}^{n+1}$ must satisfy the constraint $\zeta_{i,j}^{n+1} \in C(I_j)$, where $C(I_j) = [c_{\min}(I_j), -k_5] \cup [0, c_{\max}(I_j)]$, as defined in (2.7). Moreover, to prevent the value of $\mathcal{I}_{j(i,n+1)}^n$ from going outside of the domain $[0, I_{\max}]$, we need to impose further constraints on $\zeta_{i,j}^{n+1}$. Let $C_j^{n+1} \subseteq C(I_j)$ denote the set of values of $\zeta_{i,j}^{n+1}$ such that the resulting $\mathcal{I}_{j(i,n+1)}^n$ calculated from (3.14) is bounded within $[0, I_{\max}]$. Note that C_j^{n+1} is independent of P_i. We regard all elements in C_j^{n+1} as admissible controls.

Equation (3.14) provides $\mathcal{I}_{j(i,n+1)}^n$ as an approximation to $\mathcal{I}(\tau^n)$. Hence, $V\big(P_i, \mathcal{I}_{j(i,n+1)}^n, \tau^n\big)$ is an approximation to $V(P_i, \mathcal{I}(\tau^n), \tau^n)$, which is the value function at τ^n when P is fixed at P_i and I residing at the departure point of the path $\mathcal{I}(\tau)$. As mentioned above, $\mathcal{I}_{j(i,n+1)}^n$ usually does not coincide with a grid point in the I direction. Thus, we have to choose an interpolation scheme to approximate $V\big(P_i, \mathcal{I}_{j(i,n+1)}^n, \tau^n\big)$ using discrete grid values $V_{i,j}^n$, $i = 1, \ldots, i_{\max}$, $j = 1, \ldots, j_{\max}$. Let $V_{i,j(i,n+1)}^n$ denote an approximation of $V\big(P_i, \mathcal{I}_{j(i,n+1)}^n, \tau^n\big)$ obtained by interpolating a set of values $V_{i,j}^n$ with P fixed at P_i and I varied.

Evaluating the integrals in (3.12) at $\tau = \tau^{n+1}$ using the rectangular rule, assuming that the control path $\zeta_{i,j}(\tau)$ follows (3.13) and the semi-Lagrangian trajectory $\mathcal{I}(\tau)$ satisfies (3.14), gives

$$V_{i,j}^{n+1} = \max_{\zeta_{i,j}^{n+1} \in C_j^{n+1}} \Big\{ V_{i,j(i,n+1)}^n + \Delta\tau^n\big(\zeta_{i,j}^{n+1} - a\big(\zeta_{i,j}^{n+1}\big)\big) P_i \Big\} + \Delta\tau^n (\mathcal{L}_h V)_{i,j}^{n+1}, \tag{3.15}$$

where $V(P_i, \mathcal{I}(\tau^n), \tau^n)$ in (3.12) is approximated by $V_{i,j(i,n+1)}^n$. The last term in (3.15) follows from approximating the second integral in (3.12), assuming $\mathcal{L}V(P_i, \mathcal{I}(\tau), \tau) = \mathcal{L}V(P_i, \mathcal{I}(\tau^{n+1}), \tau^{n+1}) = \mathcal{L}V(P_i, I_j, \tau^{n+1})$ for $\tau \in [\tau^n, \tau^{n+1}]$, and then replacing the differential operator $\mathcal{L}V$ with its discrete form $\mathcal{L}_h V$, given in (3.3). Equations (3.14)–(3.15) specify a semi-Lagrangian fully implicit discretization. Assuming that the solution value is smooth, although this is not the case in general, we show in Lemma 4.5 that linear interpolation for computing $V_{i,j(i,n+1)}^n$ is sufficient to achieve a first-order global discretization error. We will also demonstrate the first-order convergence of the fully implicit timestepping scheme using numerical experiments in section 5.

3.2. Crank–Nicolson timestepping. In order to obtain a higher-order discretization in time, we can evaluate the integrals in (3.9) and (3.12) using a trapezoidal rule, which results in a Crank–Nicolson timestepping scheme. We assume that the control path $\zeta_{i,j}(\tau)$ is a continuous differentiable function of time. Let $\zeta_{i,j}^n = \zeta_{i,j}(\tau^n)$. Applying the trapezoidal rule to the integral in (3.9), assuming the control is a smooth function of time, gives the following approximation for $\mathcal{I}_{j(i,n+1)}^n$:

$$\mathcal{I}_{j(i,n+1)}^n = I_j - \frac{\Delta\tau^n}{2}\big(\zeta_{i,j}^{n+1} + a\big(\zeta_{i,j}^{n+1}\big)\big) - \frac{\Delta\tau^n}{2}\big(\zeta_{i,j}^n + a\big(\zeta_{i,j}^n\big)\big). \tag{3.16}$$

As for the definition of admissible controls in the previous subsection, we can define $C_j^{n+1} \subseteq C\big(\mathcal{I}_{j(i,n+1)}^n\big) \times C(I_j)$ to be the set of all admissible controls $\zeta_{i,j}^n$ and $\zeta_{i,j}^{n+1}$ such that the value $\mathcal{I}_{j(i,n+1)}^n$ calculated from (3.16) resides inside the domain $[0, I_{\max}]$.

Approximating the integrals in (3.12) using the trapezoidal rule and assuming that the control path $\zeta_{i,j}(\tau)$ is a smooth function of time and that the semi-Lagrangian trajectory $\mathcal{I}(\tau)$ follows (3.16), then we obtain

$$V_{i,j}^{n+1} = \max_{(\zeta_{i,j}^n, \zeta_{i,j}^{n+1}) \in C_j^{n+1}} \left\{ V_{i,j(i,n+1)}^n + \frac{\Delta\tau^n}{2}\big(\zeta_{i,j}^{n+1} - a\big(\zeta_{i,j}^{n+1}\big)\big)P_i \right.$$
$$\left. + \frac{\Delta\tau^n}{2}\big(\zeta_{i,j}^{n} - a\big(\zeta_{i,j}^{n}\big)\big)P_i + \frac{\Delta\tau^n}{2}(\mathcal{L}_h V)_{i,j(i,n+1)}^n \right\} + \frac{\Delta\tau^n}{2}(\mathcal{L}_h V)_{i,j}^{n+1}, \tag{3.17}$$

where $(\mathcal{L}_h V)_{i,j(i,n+1)}^n$ is the evaluation of the discrete differential operator (3.3) at $\tau = \tau^n$ and $(P, I) = \big(P_i, \mathcal{I}_{j(i,n+1)}^n\big)$ with $\mathcal{I}_{j(n+1)}^n$ given in (3.16). Equations (3.16)–(3.17) results in a semi-Lagrangian Crank–Nicolson discretization.

As in the fully implicit timestepping case, we need to use interpolation to compute quantities of the form $(\cdot)_{i,j(i,n+1)}^n$ in (3.17) since $\mathcal{I}_{j(n+1)}^n$ usually does not reside on a grid point in the I direction. As suggested by [17, 9, 21] for the case when the control is a fixed constant, second-order global truncation error can be achieved if the P derivatives in $\mathcal{L}V$ are discretized using second-order accurate methods, e.g., a central differencing method (see Appendix A), and quadratic interpolation used for $(\cdot)_{i,j(i,n+1)}^n$. Of course, this truncation error estimate is valid only for smooth solutions. Indeed, in the numerical experiments conducted in section 5, we cannot observe second-order convergence for the Crank–Nicolson timestepping scheme with high-order interpolants.

3.3. Matrix form of the discrete equations. Before proceeding to setting up the matrix form for the discrete equations (3.14)–(3.15) and (3.16)–(3.17), let us introduce the following notation. Let V_j^n be a column vector of discrete solution values with $[V_j^n]_i = V_{i,j}^n$ for a fixed I_j. Let V^n be a column vector with $[V^n]_j = V_j^n$. Based on the discrete differential operator $\mathcal{L}_h V$ in (3.3), we can define a matrix L^n such that

$$\begin{aligned}[L^n \cdot V^n]_{i,j} &= (\mathcal{L}_h V)_{i,j}^n \\ &= [\alpha_i^n V_{i-1,j}^n + \beta_i^n V_{i+1,j}^n - (\alpha_i^n + \beta_i^n + r)V_{i,j}^n],\end{aligned} \tag{3.18}$$

where the coefficients α_i^n and β_i^n are given in Appendix A.

Let Φ^{n+1} be a Lagrange interpolation operator such that

$$\big[\Phi^{n+1} \cdot V^n\big]_{i,j} = V_{i,j(i,n+1)}^n + \text{interpolation error}, \tag{3.19}$$

where Φ^{n+1} can represent any order (linear, quadratic) of Lagrangian interpolation. Let $\big[\Phi^{n+1}V^n\big]_j$ denote a column vector with entries

$$\Big[\big[\Phi^{n+1}V^n\big]_j\Big]_i = \big[\Phi^{n+1}V^n\big]_{i,j}. \tag{3.20}$$

Let P denote a column vector satisfying $[P]_i = P_i$. Let ζ_j^n and ζ_j^{n+1} be diagonal matrices with diagonal entries $\big[\zeta_j^n\big]_{ii} = \zeta_{i,j}^n$ and $\big[\zeta_j^{n+1}\big]_{ii} = \zeta_{i,j}^{n+1}$. Similarly, let $a\big(\zeta_j^n\big)$ and $a\big(\zeta_j^{n+1}\big)$ denote diagonal matrices with diagonal entries $\big[a\big(\zeta_j^n\big)\big]_{ii} = a\big(\zeta_{i,j}^n\big)$, $\big[a\big(\zeta_j^{n+1}\big)\big]_{ii} = a\big(\zeta_{i,j}^{n+1}\big)$. Let I be an identity matrix. Given the above notation, the discrete equations (3.14)–(3.15) and (3.16)–(3.17) together can be written in a θ-form

as

$$\begin{aligned}\big[[I-(1-\theta)\Delta\tau^n L^{n+1}]\,V^{\,n+1}\big]_j &= \big[\Phi^{n+1}\big[I+\theta\Delta\tau^n L^n\big]V^n\big]_j \\ &\quad + (1-\theta)\Delta\tau^n\big(\zeta_j^{n+1}-a\big(\zeta_j^{n+1}\big)\big)P+\theta\Delta\tau^n\big(\zeta_j^n-a\big(\zeta_j^n\big)\big)P,\end{aligned} \tag{3.21}$$

$$\begin{aligned}\text{where } \big[\zeta_j^n\big]_{ii},\big[\zeta_j^{n+1}\big]_{ii} = \underset{([\zeta_j^n]_{ii},[\zeta_j^{n+1}]_{ii})\in C_j^{n+1}}{\operatorname{argmax}} \Big\{\Big[\big[\Phi^{n+1}[I+\theta\Delta\tau^n L^n]V^n\big]_j \\ + (1-\theta)\Delta\tau^n\big(\zeta_j^{n+1}-a\big(\zeta_j^{n+1}\big)\big)P+\theta\Delta\tau^n\big(\zeta_j^n-a\big(\zeta_j^n\big)\big)P\Big]_i\Big\}\end{aligned}$$

for $j=1,\dots,j_{\max}$. Here $\theta=0$ corresponds to fully implicit timestepping, and $\theta=1/2$ is Crank–Nicolson timestepping.

It will be convenient to define $\Delta P_{\max}=\max_i\big(P_{i+1}-P_i\big)$, $\Delta P_{\min}=\min_i\big(P_{i+1}-P_i\big)$, $\Delta I_{\max}=\max_j\big(I_{j+1}-I_j\big)$, $\Delta I_{\min}=\min_j\big(I_{j+1}-I_j\big)$, $\Delta\tau_{\max}=\max_n\big(\tau^{n+1}-\tau^n\big)$, $\Delta\tau_{\min}=\min_n\big(\tau^{n+1}-\tau^n\big)$. We assume that there are mesh size/timestep parameters $h_{\min}=C_0h_{\max}$ such that

$$\begin{aligned}&\Delta P_{\max}=C_1h_{\max},\quad \Delta I_{\max}=C_2h_{\max},\quad \Delta\tau_{\max}=C_3h_{\max},\\ &\Delta P_{\min}=C_1'h_{\min},\quad \Delta I_{\min}=C_2'h_{\min},\quad \Delta\tau_{\min}=C_3'h_{\min},\end{aligned} \tag{3.22}$$

where $C_0,C_1,C_1',C_2,C_2',C_3,C_3'$ are constants independent of $h_{\min},h_{\max}$. We can write the discrete equations (3.21) at each node (P_i,I_j) as

(3.23)

$$\begin{aligned}&G_{i,j}^{n+1}\big(h_{\max},V_{i,j}^{n+1},\big\{V_{k,j}^{n+1}\big\}_{k\neq i},\big\{V_{i,j}^n\big\}\big)\\ &\equiv \min_{(\zeta_{i,j}^n,\zeta_{i,j}^{n+1})\in C_j^{n+1}}\Bigg\{\frac{V_{i,j}^{n+1}-\big[\Phi^{n+1}V^n\big]_{i,j}}{\Delta\tau^n}-(1-\theta)\big[L^{n+1}V^{n+1}\big]_{i,j}-\theta\big[\Phi^{n+1}L^nV^n\big]_{i,j}\\ &\qquad -(1-\theta)\big[\big(\zeta_j^{n+1}-a\big(\zeta_j^{n+1}\big)\big)P\big]_i-\theta\big[\big(\zeta_j^n-a\big(\zeta_j^n\big)\big)P\big]_i\Bigg\}\\ &=0,\end{aligned}$$

where $\big\{V_{k,j}^{n+1}\big\}_{k\neq i}$ is the set of values $V_{k,j}^{n+1}$, $k\neq i$, $k=1,\dots,i_{\max}$, and $\big\{V_{i,j}^n\big\}$ is the set of values $V_{i,j}^n$, $i=1,\dots,i_{\max}$, $j=1,\dots,j_{\max}$.

4. Convergence to the viscosity solution. The pricing PDE (2.10) is nonlinear; hence the solution to the PDE may not be unique [14, 2]. As pointed out in [23, 2], it is important to ensure that a numerical scheme for solving PDE (2.10) converges to the valid solution from a financial perspective, which in this case is the viscosity solution. Provided that a strong comparison result for the PDE applies, [7, 2] demonstrate that a numerical scheme will converge to the viscosity solution of the equation if it is l_∞ stable, pointwise consistent, and monotone. Schemes failing to satisfy these conditions may converge to nonviscosity solutions. In fact, [29] gives an example where seemingly reasonable discretizations of nonlinear option pricing PDEs that do not satisfy the sufficient convergence conditions for viscosity solutions either never converge or converge to a nonviscosity solution.

We first discuss the existence of the strong comparison result for (2.10).

4.1. Strong comparison result. There are various research papers deriving a strong comparison result for second-order HJB equations associated with several types of boundary conditions [4, 6, 12, 3, 10]. In particular, [4, 6] prove that the viscosity solution of degenerate elliptic HJB equations with Dirichlet boundary conditions satisfies the strong comparison result, provided that several assumptions are satisfied. In [6], the author demonstrates that the following hold:

S1. If the coefficient of the diffusion term (in our case $(\hat{\sigma}(P)P)^2$) vanishes at a region on a boundary with an outgoing characteristic, independent of the value for the control variable, then the viscosity solution on this boundary region is the limit of the viscosity solution from interior points.

S2. If the characteristic at a region on the boundary, associated with the first-order term in the PDE, is incoming to the domain independent of the choice of the control value, then the viscosity solution at the region corresponds to the specified boundary data in the classical sense.

We can regard the two-dimensional parabolic PDE (2.10) as a three-dimensional degenerate elliptic PDE in the variable $x = (\tau, P, I) \in \mathbb{R}^N \times \mathbb{R}^N \times [0, I_{\max}]$. The resulting elliptic PDE is degenerate in the sense that the equation does not contain the second-order derivatives with respect to τ and I, or, equivalently, the effective volatility (i.e., the diffusion term) is zero with respect to τ and I. We solve PDE (2.10) in the boundary region $(P, I, \tau) \in [0, P_{\max}] \times \{0, I_{\max}\} \times (0, T)$. Conditions (2.12)–(2.13) imply that the statement S1 above is satisfied for this boundary region. In the boundary region $(P, I, \tau) \in \{0, P_{\max}\} \times [0, I_{\max}] \times (0, T)$ we solve (2.14)–(2.15). Since $\alpha(K(t) - P) \geq 0$ as $P \to 0$ and $\alpha(K(t) - P) \leq 0$ as $P \to P_{\max}$, statement S1 above is also satisfied for this region. Thus, the viscosity solution does not require boundary data in both the P and I directions, which confirms our intuition in setting the boundary conditions in these directions. PDE (2.10) implies that statement S2 above is satisfied in the region when $(P, I, \tau) \in [0, P_{\max}] \times [0, I_{\max}] \times \{0\}$. This means that the viscosity solution uses the Dirichlet boundary condition, which we provided as the payoff function in (2.11).

From the analysis above, the boundary conditions we apply for (2.10) are in accordance with the behavior of the viscosity solution at the boundary. Consequently, we can use the strong comparison result in [4, 6] if (2.10)–(2.15) satisfy assumptions given in [4, 6]. However, a technical difficulty arises when we try to verify an assumption among those outlined in [4, 6]: the boundary is assumed to be smooth in [4, 6], so that the distance function from a point in the interior to the boundary is well defined. In our case, however, the boundary surface is a cuboid, which results in nonsmoothness of the distance function in the corners of the cuboid. In [12], the strong comparison result is proved for a similar (but not identical) problem associated with a nonsmooth boundary. Consequently, we make the following assumption, which is necessary to ensure that a unique viscosity solution to (2.10) exists.

Assumption 4.1. Pricing equation (2.10) and the associated boundary conditions (2.11)–(2.15) satisfy the strong comparison result as defined in [7, 2].

Based on Assumption 4.1, we will show that the semi-Lagrangian fully implicit discretization (3.21) or (3.23) (with $\theta = 0$) converges to the viscosity solution of (2.10) by verifying the stability, consistency, and monotonicity of the scheme in sequence. As explained in section 3.2, higher than or equal to third-order (quadratic) interpolation is needed for operation Φ^{n+1} in (3.23) in order to achieve a second-order global truncation error for Crank–Nicolson timestepping (for smooth solutions). This makes this scheme nonmonotone in general, and hence we cannot guarantee convergence

of high-order Crank–Nicolson timestepping to the viscosity solution because monotonicity can be obtained only for linear interpolation. We will, nevertheless, prove the consistency of the Crank–Nicolson timestepping scheme and carry out numerical experiments with Crank–Nicolson timestepping using quadratic interpolation.

4.2. Stability. In this section, we demonstrate the l_∞ stability of the scheme (3.21).

DEFINITION 4.2 (stability). *Discretization* (3.21) *is* l_∞ *stable if*

$$\|V^{n+1}\|_\infty \le C_4 \tag{4.1}$$

for $0 \le n \le N$ *as* $\Delta\tau_{\min} \to 0$, $\Delta P_{\min} \to 0$, $\Delta I_{\min} \to 0$, *where* C_4 *is a constant independent of* $\Delta\tau_{\min}$, $\Delta P_{\min}$, $\Delta I_{\min}$. *Here* $\|V^{n+1}\|_\infty = \max_{i,j} |V^{n+1}_{i,j}|$.

The stability of the semi-Lagrangian fully implicit discretization (3.21) is a consequence of the following lemma.

LEMMA 4.3 (stability of the fully implicit scheme). *Assuming that discretization* (3.3) *satisfies the positive coefficient condition* (3.4) *and that linear interpolation is used to compute* $V^n_{i,j(i,n+1)}$, *then the scheme* (3.21) *satisfies*

$$\|V^{n+1}\|_\infty \le \|V^0\|_\infty + C_5 \tag{4.2}$$

in the case of fully implicit timestepping ($\theta = 0$), *where*

$$C_5 = T \cdot P_{\max} \cdot \max\{|c_{\max}(I_{\max})|, |c_{\min}(0)|\}.$$

Proof. The proof directly follows from applying the maximum principle to the discrete equation (3.21). We omit the details here. Readers can refer to [17, Theorem 5.5] and [23] for complete stability proofs of the semi-Lagrangian fully implicit scheme for American Asian options and that of finite difference schemes for controlled HJB equations, respectively. ∎

4.3. Consistency. We give in the following a definition for the consistency of a discretization.

DEFINITION 4.4 (consistency). *The scheme* $G^{n+1}_{i,j}\big(h_{\max}, V^{n+1}_{i,j}, \{V^{n+1}_{k,j}\}_{k\ne i}, \{V^n_{i,j}\}\big)$ *given in* (3.23) *is consistent if, for any smooth test function* $\phi(P, I, \tau)$ *with bounded derivatives of all orders, with* $\phi^n_{i,j} = \phi(P_i, I_j, \tau^n)$, *we have that*

$$\begin{aligned}\lim_{h_{\max}\to 0} \Bigg| &\min_{c\in C(I_j)} \Big\{ \big(\phi_\tau + (c + a(c))\phi_I - (c - a(c))P - \mathcal{L}\phi\big)^{n+1}_{i,j} \Big\} \\ &- G^{n+1}_{i,j}\big(h_{\max}, \phi^{n+1}_{i,j}, \{\phi^{n+1}_{k,j}\}_{k\ne i}, \{\phi^n_{i,j}\}\big) \Bigg| = 0.\end{aligned} \tag{4.3}$$

Note that in order to handle the boundary data, the consistency condition is more complicated if general Dirichlet or Neumann boundary conditions are specified [7, 2]. In our case, the boundary conditions specified in section 2.4 are the limit of (2.10) as P, I go towards the boundary. As a result, the need for this more general condition does not occur.

LEMMA 4.5 (consistency). *Suppose that the mesh size and timestep size satisfy* (3.22), *and that the control parameters satisfy condition* (2.7). *Then the discretization* (3.23) *is consistent as defined in Definition* 4.4, *provided that* $c_{\max}$, $c_{\min}$, *and* $a(c)$ *satisfy* (2.1), (2.3), *and* (2.5), *respectively. In particular, assuming that linear*

interpolation is used in operation Φ^{n+1} *in* (3.19), *the global discretization error of the scheme* $G_{i,j}^{n+1}$ *is* $O(h_{\max})$.

Proof. We provide only an outline of the proof in Appendix B. Refer to [13, Appendix D] for a detailed proof. ∎

4.4. Monotonicity. In this section, we discuss the monotonicity of the fully implicit scheme (3.23).

DEFINITION 4.6 (monotonicity). *The scheme* $G_{i,j}^{n+1}$ $(h_{\max},\ V_{i,j}^{n+1},\ \{V_{k,j}^{n+1}\}_{k\neq i},\ \{V_{i,j}^{n}\})$ *given in* (3.23) *is monotone if*

$$\begin{aligned} &G_{i,j}^{n+1}\big(h_{\max}, V_{i,j}^{n+1}, \big\{V_{k,j}^{n+1}+\epsilon_{k,j}^{n+1}\big\}_{k\neq i}, \big\{V_{i,j}^{n}+\epsilon_{i,j}^{n}\big\}\big) \\ &\quad\le G_{i,j}^{n+1}\big(h_{\max}, V_{i,j}^{n+1}, \big\{V_{k,j}^{n+1}\big\}_{k\neq i}, \big\{V_{i,j}^{n}\big\}\big) \quad \forall i,j,\ \forall \epsilon_{i,j}^{n}, \epsilon_{i,j}^{n+1}\ge 0, \end{aligned} \tag{4.4}$$

and, for any constants $m, m_0 \ge 0$, *we have*

$$\begin{aligned} &G_{i,j}^{n+1}\big(h_{\max}, V_{i,j}^{n+1}+m+m_0\tau^{n+1}, \big\{V_{k,j}^{n+1}+m+m_0\tau^{n+1}\big\}_{k\neq i}, \big\{V_{i,j}^{n}+m+m_0\tau^{n}\big\}\big) \\ &\quad\ge m_0 + G_{i,j}^{n+1}\big(h_{\max}, V_{i,j}^{n+1}, \big\{V_{k,j}^{n+1}\big\}_{k\neq i}, \big\{V_{i,j}^{n}\big\}\big)\ \forall i,j,\ \forall m, m_0\ge 0. \end{aligned} \tag{4.5}$$

This definition of monotonicity is equivalent to that introduced in [24].

LEMMA 4.7 (monotonicity). *If the discretization* (3.3) *satisfies the positive coefficient condition* (3.4) *and linear interpolation is used to compute* $V_{i,j(i,n+1)}^{n}$, *then in the case of fully implicit timestepping* ($\theta = 0$), *the discretization* $G_{i,j}^{n+1}$ $(h_{\max},\ V_{i,j}^{n+1},\ \{V_{k,j}^{n+1}\}_{k\neq i},\ \{V_{i,j}^{n}\})$, *as given in* (3.23), *is monotone according to Definition* 4.6.

Proof. The proof directly follows that of monotonicity of finite difference schemes for controlled HJB equations in [5, 23]. Note that the proof assumes the consistency of the discretization (3.23), which has been shown in Lemma 4.5. ∎

From Assumption 4.1 and Lemmas 4.3, 4.5, and 4.7, and using the results in [7, 2], we can obtain the following convergence result.

THEOREM 4.8 (convergence to the viscosity solution). *Assuming that discretization* (3.21) *satisfies Assumption* 4.1 *and all the conditions required for Lemmas* 4.3, 4.5, *and* 4.7, *then scheme* (3.21) *converges to the viscosity solution of* (2.10) *in the case of fully implicit timestepping* ($\theta = 0$).

4.5. Solution algorithm. The solution to (2.10) can be computed using Algorithm 4.9.

Remark 4.1. As described in [23], a standard implicit finite difference discretization for (2.10) requires a Policy-type iteration at each timestep to solve the nonlinear discretized algebraic equations. An alternative approach uses an explicit timestepping method, but an explicit method suffers from the usual parabolic stability condition. However, Algorithm 4.9, which uses a semi-Lagrangian discretization, avoids the need for Policy iteration. Instead, Algorithm 4.9 replaces the nonlinearity involving V^{n+1} with local optimization problems involving V^n, the known solution from the previous timestep.

Remark 4.2. At each timestep in Algorithm 4.9 all discrete equations in the I direction are decoupled and independent of each other. This property makes solution of the gas storage contract straightforward to implement.

ALGORITHM 4.9. Semi-Lagrangian timestepping.

$V^0 =$ Option Payoff
For $n = 0, \ldots,$ // Timestep loop
  For $j = 1, \ldots,$ // Loop through nodes in I direction
    For $i = 1, \ldots,$ // Loop through nodes in P direction

$$\left[\zeta_j^n\right]_{ii}, \left[\zeta_j^{n+1}\right]_{ii} = \underset{([\zeta_j^n]_{ii}, [\zeta_j^{n+1}]_{ii}) \in C_j^{n+1}}{\operatorname{argmax}} \left\{ \left[\left[\Phi^{n+1}[I + \theta\Delta\tau^n L^n]V^n\right]_j + (1-\theta)\Delta\tau^n\left(\zeta_j^{n+1} - a\left(\zeta_j^{n+1}\right)\right)P + \theta\Delta\tau^n\left(\zeta_j^n - a\left(\zeta_j^n\right)\right)P \right]_i \right\} \tag{4.6}$$

    EndFor
    Solve

$$\left[[I - (1-\theta)\Delta\tau^n L^{n+1}]V^{n+1}\right]_j = \left[\Phi^{n+1}\left[I + \theta\Delta\tau^n L^n\right]V^n\right]_j + (1-\theta)\Delta\tau^n\left(\zeta_j^{n+1} - a\left(\zeta_j^{n+1}\right)\right)P + \theta\Delta\tau^n\left(\zeta_j^n - a\left(\zeta_j^n\right)\right)P$$

  EndFor
EndFor

Note that for a finite grid size, the solution of the discrete control problem in Algorithm 4.9 may allow controls which are not possible controls for the exact solution of the HJB equation. The exact solution of (2.10) has the property that the controls are of the *bang-bang* type [31]; i.e., the optimal controls can take on only the values in a finite set. Consequently, there are two possible approaches for determining the optimal controls at each grid node. First, we can use our knowledge of the exact controls to search only for optimal controls within the known finite set of possible values. In other words, the set of admissible controls is finite in this case; hence this approach is consistent with the control behavior in the exact solution. We refer to this approach as the *bang-bang* method. Alternatively, we can simply solve the discrete optimization problem in Algorithm 4.9 and allow any admissible control. We will refer to this technique as the *no bang-bang* method.

4.6. Computational details. Recall that in Algorithm 4.9 we need to solve a local optimization problem for each node (P_i, I_j, τ^n). In [13, Appendix C] we provide the details of two methods for carrying out this operation.

In the no bang-bang case, if linear interpolation is used in (3.19), then the maximum can also be determined by examining only a finite number of nodal values (for each (P_i, I_j, τ^n)). Since linear interpolation is correct to $O(h_{\max}^2)$ for any smooth function, then the maximum determined with this approach is consistent with the exact maximum for any smooth test function.

Since it is known that the solution to PDE (2.10) has bang-bang-type controls and since the no bang-bang method converges to the viscosity solution, both methods will converge to the solution with bang-bang-type controls, assuming that a fully implicit timestepping scheme is used.

It is easy to see that, from (3.5), (3.16), and (3.22), the number of nodes that must be examined to solve the local optimization problem at each node, using the above methods, is a constant independent of $h_{\min}$. Since (3.5) is independent of P, we can precompute and store interpolation indices and weights. In addition, the linear system solve has been reduced to a set of decoupled one-dimensional problems. Hence

the complexity per timestep of the implicit semi-Lagrangian scheme is linear in the total number of nodes. Thus the complexity per timestep of the fully implicit semi-Lagrangian scheme is the same as an explicit method, but has the obvious advantage of being unconditionally stable.

5. Numerical experiments. Having presented several semi-Lagrangian discretization schemes in the previous sections, in this section we conduct numerical experiments based on these schemes. We use "dollars per million British thermal units" ($/mmBtu) and "million cubic feet" (MMcf) as the default units for gas spot price P and gas inventory I, respectively. Since 1000 mmBtus are roughly equal to 1 MMcf, in order to unify the units, we need to multiply gas spot price by 1000 when computing payoffs or revenues.

Throughout the numerical experiments, we use the following nonsmooth payoff function from [11]:

$$V(P, I, t = T) = -2P \max(1000 - I, 0). \tag{5.1}$$

Equation (5.1) indicates that severe penalties are charged if the gas inventory is less than 1000 MMcf, and no compensation is received when the inventory is above 1000 MMcf. Naturally, such a payoff structure will force the operator of a gas storage facility to maintain the gas inventory as close to 1000 MMcf as possible at maturity to avoid revenue loss.

This section is arranged as follows: we first give numerical results for the case without incorporating the seasonality effect into the equilibrium natural gas spot price; we then incorporate the seasonality feature and illustrate its influence on both the solution value and the optimal control strategy. At the end of this section, we further extend the underlying risk-adjusted gas spot price process to include a compound Poisson process that simulates random jumps of the gas prices, and then present numerical results incorporating a jump diffusion process.

5.1. No seasonality effect. In this subsection, we assume that the equilibrium gas price is independent of time; that is, we set $K(t) = K_0$ in (2.9).

We first carry out a convergence analysis, assuming that the risk-adjusted natural gas spot price follows the pure mean-reverting process (2.8) with $\alpha = 2.38$, $K_0 = 6$, $\sigma = 0.59$. In other words, the risk-adjusted gas spot price follows

$$dP = 2.38(6 - P)dt + 0.59PdZ. \tag{5.2}$$

We are most interested in the solution when the gas spot price is near the long-term equilibrium price, i.e., $P = 6$ $/mmBtu for (5.2). Note that when $I = 1000$ MMcf, the payoff is nonsmooth (see (5.1)). Consequently, to fully test our semi-Lagrangian discretization schemes, we focus on the convergence results at $(P, I) = (6, 1000)$. We use an unequally spaced grid in the P, I directions, where there are more nodes around the mesh point $(P, I) = (6, 1000)$ compared with other locations.

Table 5.1 lists other input parameters for pricing the value of a gas storage contract. The convergence results obtained from refining the mesh spacing and timestep size are shown in Table 5.2, where we use fully implicit and Crank–Nicolson timestepping schemes associated with both the bang-bang and no bang-bang methods for solving the discrete optimization problem in Algorithm 4.9. Linear interpolation and quadratic interpolation are used for fully implicit and Crank–Nicolson timestepping, respectively. (See section 3 for a discussion on interpolation schemes.) Following [29], in order to improve the convergence for nonsmooth payoff (5.1), we use a modification

TABLE 5.1

Input parameters used to price the value of a gas storage contract, where $I_{\max}$ is the maximum storage inventory; k_1, k_2, k_3, k_4, k_5 are parameters in (2.1)–(2.2) *and* (2.5). *The values of $I_{\max}, k_1, k_2, k_3, k_4, k_5$ are taken from* [31].

Parameter	Value	Parameter	Value
r	0.1	k_2	730000
T	3 years	k_3	500
$I_{\max}$	2000 MMcf	k_4	2500
k_1	2040.41	k_5	$1.7 \cdot 365$

TABLE 5.2

Value of a natural gas storage facility at $P = 6$ \$/mmBtu and $I = 1000$ MMcf. Risk-adjusted gas spot price follows the pure mean-reverting process (5.2). *Convergence ratios are presented for fully implicit and Crank–Nicolson timestepping schemes with the bang-bang and the no bang-bang methods. Constant timesteps are used. Payoff function is given in* (5.1). *Other input parameters are given in Table* 5.1. *Crank–Nicolson incorporates the modification suggested in* [30].

P grid nodes	I grid nodes	No. of timesteps	Bang-bang method		No bang-bang method	
			Value	Ratio	Value	Ratio
Fully implicit timestepping						
53	61	500	4477036	n.a.	4556380	n.a.
105	121	1000	4503705	n.a.	4542845	n.a.
209	241	2000	4514723	2.42	4534660	1.63
417	481	4000	4519809	2.17	4530331	1.89
833	961	8000	4522653	1.79	4528219	2.05
Crank–Nicolson timestepping						
53	61	500	4483667	n.a.	4520475	n.a.
105	121	1000	4509076	n.a.	4525352	n.a.
209	241	2000	4517960	2.86	4526280	5.26
417	481	4000	4522632	1.90	4527225	0.98
833	961	8000	4524948	2.02	4527331	8.92

suggested by [30] for Crank–Nicolson timestepping. Specifically, we apply fully implicit timestepping in the first four timesteps, and use Crank–Nicolson timestepping in the rest of the timesteps.

The results in Table 5.2 indicate that both timestepping schemes converge to the same solution, although convergence to the viscosity solution can only be guaranteed for fully implicit timestepping given Assumption 4.1. We define the convergence ratio as the ratio of successive changes in the solution, as the timestep and mesh size are reduced by a factor of two. A ratio of two indicates first-order convergence, while a ratio of four indicates second-order convergence. The convergence ratios are approximately two for fully implicit timestepping with both the bang-bang and the no bang-bang methods. Note that the no bang-bang method is a more general approach which can be used in cases where controls are not of bang-bang type.

It is interesting to note that in a fixed refinement level, the bang-bang method results in a smaller value than the no bang-bang method for both timestepping schemes. This is because the no bang-bang method actually solves the discrete optimization problem in Algorithm 4.9, instead of only testing a finite set of points, which results in a higher solution value for PDE (2.10) (for a finite grid size) than the bang-bang method.

For the fully implicit tests in Table 5.2, the no bang-bang method requires about 10% more CPU time than does the bang-bang method. This is consistent with our earlier estimates, since the the no bang-bang examines only a constant number of grid nodes per node in order to solve the local optimization problem.

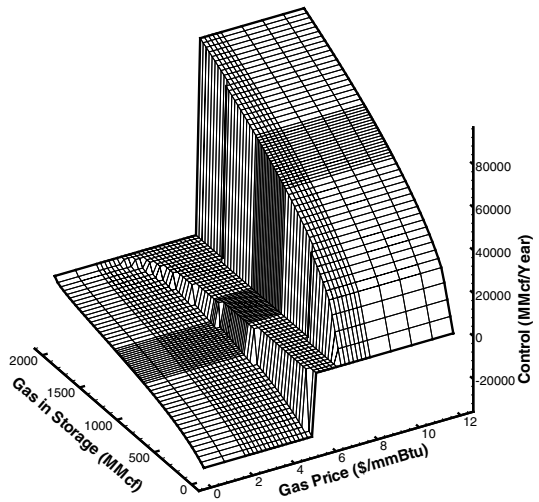

FIG. 5.1. *Optimal control strategy at current time $t = 0$ as a function of gas spot price P and gas inventory I. Risk-adjusted gas spot price follows the pure mean-reverting process* (5.2). *Payoff function is given in* (5.1). *Other input parameters are given in Table* 5.1. *Fully implicit timestepping with the no bang-bang method and with constant timesteps is used.*

Table 5.2 also shows that Crank–Nicolson timestepping does not appear to converge at a second-order rate. We have observed this same effect in many of our tests. Since we do not seem to obtain any benefit from Crank–Nicolson timestepping, fully implicit timestepping appears to be a better choice since we are guaranteed convergence to the viscosity solution given Assumption 4.1, as shown in section 4. In the rest of this paper, we will use fully implicit timestepping exclusively.

Figure 5.1 shows the optimal control strategy surface at $t = 0$ as a function of P and I. This surface is similar to that given in [31]. The interpretation given in [31] also applies to Figure 5.1.

Our numerical computations truncate the domain $P \in [0, \infty]$ to $[0, P_{\max}]$. In order to test the influence of the domain truncation on the solution, we compute the solution values at $P = 6$ \$/mmBtu, $I = 1000$ MMcf using two different values of $P_{\max}$: $P_{\max} = 2000$ and 20000 \$/mmBtu. We found that for all four refinement levels, the first ten digits of the two solution values are identical. This indicates that by setting $P_{\max} = 2000$ \$/mmBtus, there is a negligible solution error incurred by the domain truncation. As a result, all subsequent results will be reported using $P_{\max} = 2000$ \$/mmBtus.

If an explicit, semi-Lagrangian scheme is used, then the stability condition is

$$\Delta\tau^n < \min_i \left\{ \frac{1}{\alpha_i^n + \beta_i^n + r} \right\}, \tag{5.3}$$

where parameters α_i^n, β_i^n are given in Appendix A. Condition (5.3) implies that $\Delta\tau^n = O((\Delta P_{\min})^2)$, where $\Delta P_{\min} = \min_i(P_{i+1} - P_i)$. In contrast, there is no such timestep restriction for fully implicit timestepping.

In [31], a fully explicit method (standard timestepping) is used for the gas storage

TABLE 5.3

Value of a natural gas storage facility at $P = 6$ \$/mmBtu and $I = 1000$ MMcf. Risk-adjusted gas spot price follows the pure mean-reverting process (5.4) *that incorporates the seasonality effect. Convergence ratios are presented for fully implicit timestepping with the bang-bang and the no bang-bang methods. Constant timesteps are used. Payoff function is given in* (5.1). *Other input parameters are given in Table* 5.1.

P grid nodes	I grid nodes	No. of timesteps	Bang-bang method		No bang-bang method	
			Value	Ratio	Value	Ratio
53	61	500	4815891	n.a.	4889602	n.a.
105	121	1000	4839796	n.a.	4875449	n.a.
209	241	2000	4848843	2.64	4866953	1.67
417	481	4000	4853100	2.13	4862530	1.92
833	961	8000	4855485	1.78	4860397	2.07

problem. In this case, the stability condition would be $\Delta\tau^n = O((\Delta P_{\min})^2 + \Delta I_{\min})$, where $\Delta I_{\min} = \min_j(I_{j+1} - I_j)$.

Recall from section 4.6 that both the implicit method used here and an explicit method have complexity linear in the number of space nodes per timestep. In general, the constant in the complexity estimate will favor the explicit standard timestepping scheme [31], due to the the extra interpolation operations required by the implicit method. In terms of running time, we expect that the implicit method will be superior to the explicit method if the spatial error dominates, since the explicit stability condition will force smaller timesteps than is required for accuracy.

On the other hand, there will undoubtedly also be cases where the error is dominated by the timestepping error, in which case an explicit method may require less running time.

We remark that the fully implicit method has the practical advantage that we are completely free to place (P, I) nodes wherever is deemed necessary, since this has no effect on the permitted timestep size.

5.2. Incorporating the seasonality effect. In this subsection, we present numerical results after incorporating the seasonality effect into the equilibrium price of the pure mean-reverting process (5.2). We modified process (5.2) to

$$dP = 2.38(6 + \sin(4\pi t) - P)dt + 0.59PdZ, \tag{5.4}$$

where the additional term $\sin(4\pi t)$ makes the equilibrium price a periodic function to represent summer and winter peaks in the equilibrium price. The convergence results for this case are shown in Table 5.3. Comparing Table 5.3 with Table 5.2 indicates that incorporating the seasonality component does not affect the convergence ratio but does increase the solution value for a fixed refinement level. This is reasonable, since the seasonality effect gives the operator of a gas storage facility an opportunity for obtaining greater profits by using an optimal control strategy that takes advantage of the seasonality feature. For example, a simple strategy of buying and storing gas in spring and then producing and selling gas in summer can normally produce profits from the seasonality effect.

Figure 5.2 shows the optimal control strategy in the seasonality case that evolves over time as a function of P when the inventory is fixed at $I = 1000$ MMcf. The figure suggests that the optimal strategy is to inject gas at the maximum rate (corresponding to the negative control region in the surface) when the gas price is low, to produce gas at the maximum rate (corresponding to the positive control region) when the gas price is high, and to do nothing (corresponding to the zero control region) when the gas

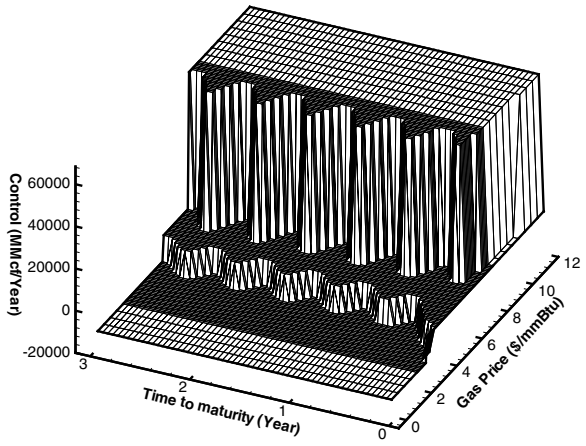

FIG. 5.2. *Optimal control strategy as a function of time-to-maturity $\tau = T - t$ and gas spot price P when the gas inventory resides at $I = 1000$ MMcf. Risk-adjusted gas spot price follows the pure mean-reverting process* (5.4)*, with seasonality. The payoff function is given in* (5.1)*. Other input parameters are given in Table* 5.1*. Fully implicit timestepping with the no bang-bang method and with constant timesteps is used.*

price is near the long-term equilibrium price. From the figure, we can clearly notice the effect of the seasonality on the control strategy: the boundary curve between the zero control region and the negative/positive control region, which represents the control switching boundary between no operation and injecting/producing gas, is periodic when it is far from maturity. We can also observe that when the contract is close to maturity, the zero control region expands rapidly. This phenomenon is caused by the payoff function (5.1): at $I = 1000$ MMcf, when close to maturity, the operator tends to stop producing gas to avoid the severe penalty at maturity. In addition, the operator will stop injecting, since any leftover gas is lost to the operator.

To illustrate the difference of the optimal control strategies before and after incorporating the seasonality effect, Figure 5.3 shows the control switching boundary curves at $I = 1000$ MMcf as a function of time-to-maturity with respect to processes (5.2) and (5.4).

5.3. Incorporating the jump effect. It is not uncommon to see spot gas price jumps, when gas is used to power electrical generating plants in times of high electricity demand. Spot gas price can jump by as much as 20% in a single day. To model this effect, in this subsection, we take the pure mean-reverting process (2.8) and extend it to include a compound Poisson process representing the jump effect, and present numerical results including a jump diffusion process. After adding a jump component, process (2.8) becomes

$$dP = [\alpha(K(t) - P) - \lambda\kappa P]dt + \sigma P dZ + (\eta - 1)P dq, \tag{5.5}$$

where the following hold:

- dq is the independent Poisson process $= \begin{cases} 0 & \text{with probability } 1 - \lambda dt, \\ 1 & \text{with probability } \lambda dt. \end{cases}$

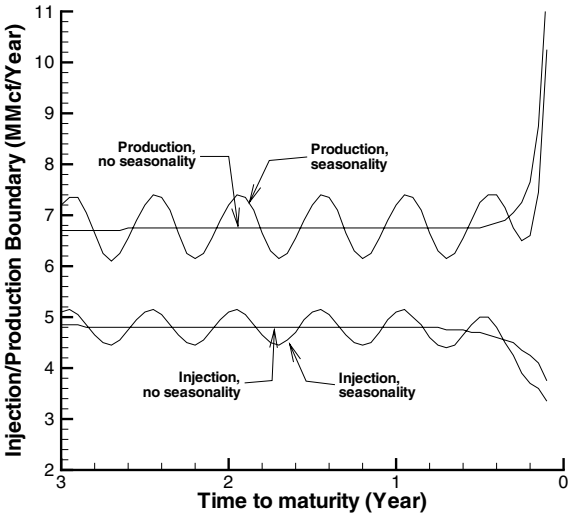

FIG. 5.3. *Control switching boundary curves as a function of time-to-maturity* $\tau = T - t$ *with respect to processes* (5.2) *(without incorporating the seasonality effect) and* (5.4) *(incorporating the seasonality effect) at* $I = 1000$ *MMcf. Payoff function is given in* (5.1). *Other input parameters are given in Table* 5.1. *Fully implicit timestepping with the no bang-bang method and with constant timesteps is used.*

- λ is the jump intensity representing the mean arrival time of the Poisson process.
- η is a random variable representing the jump size of gas price—when $dq = 1$, price jumps from P to $P\eta$. We assume that η follows a probability density function $g(\eta)$.
- κ is $E[\eta - 1]$, where $E[\,\cdot\,]$ is the expectation operator.

Assuming that the risk-adjusted gas spot price follows the jump diffusion process (5.5), the pricing PDE (2.10) turns into the following controlled partial integrodifferential equation (PIDE):

$$V_\tau = \frac{1}{2}\sigma^2 P^2 V_{PP} + [\alpha(K(t) - P) - \lambda\kappa P]V_P + \max_{c\in C(I)}\{(c - a(c))P - (c + a(c))V_I\}$$

$$(5.6)\quad -rV + \left(\lambda \int_0^\infty V(P\eta)g(\eta)d\eta - \lambda V\right).$$

Since there is no control variable in the integral term of PIDE (5.6), we can use the methods described in [17, 15, 16] to extend the semi-Lagrangian discretization schemes introduced in previous sections to solve the PIDE without difficulty. We note that it is straightforward to combine the methods in this paper with the approaches in [17] to show that the resulting scheme is consistent, stable, and monotone.

During our numerical experiments, we assume that the probability density function $g(\eta)$ follows a log-normal distribution,

$$(5.7)\qquad g(\eta) = \frac{1}{\sqrt{2\pi}\gamma\eta}\exp\left(-\frac{(\log(\eta) - \nu)^2}{2\gamma^2}\right),$$

TABLE 5.4

Input parameters for the jump diffusion process (5.5) *and the log-normal density function* (5.7). *The parameters of the jump size density function are selected so that* $E[(\eta-1)] = 0$ *and* $E[(\eta-1)^2] = .04$.

Parameter	Value	Parameter	Value
α	2.38	ν	-0.0196
$K(t)$	$6+\sin(4\pi t)$	γ	0.198
σ	0.59	λ	12

TABLE 5.5

Value of a natural gas storage facility at $P = 6$ *\$/mmBtu and* $I = 1000$ *MMcf. Risk-adjusted gas spot price follows the pure mean-reverting process* (5.5) *(incorporating the seasonality and the jump effects). Convergence ratios are presented for fully implicit timestepping with the bang-bang and the no bang-bang methods. Constant timesteps are used. Payoff function is given in* (5.1). *Input parameters are given in Tables* 5.4 *and* 5.1.

P grid nodes	I grid nodes	No. of timesteps	Bang-bang method		No-bang-bang method	
			Value	Ratio	Value	Ratio
79	61	500	7995143	n.a.	8070698	n.a.
157	121	1000	7962386	n.a.	7999775	n.a.
313	241	2000	7951062	2.89	7971737	2.53
625	481	4000	7951032	377	7961554	2.75
1249	961	8000	7951976	-0.03	7957509	2.52

with expectation $E[\eta] = \exp(\nu + \gamma^2/2)$. We will choose values of parameters ν and γ such that $(\eta-1)$, the relative change in the gas spot price, has mean zero and variance 0.04.

Table 5.4 lists values of the parameters for process (5.5) and for the log-normal density function (5.7), where the parameters of the drift and diffusion components in process (5.5) take the same values as those in process (5.4) for the case without incorporating the jump effect. Note that we set the jump intensity $\lambda = 12$ so that random jumps appear approximately once every month. Table 5.5 presents the convergence results for the solution to PIDE (5.6). Table 5.5 shows that the jump effect greatly increases the value of the storage facility.

Since the controls in (5.6) do not appear in the integral terms, it seems reasonable to suppose that the converged controls for problem (5.6) will also be of the bang-bang type, but we are not aware of a proof of this. We will solve (5.6) using both bang-bang and no bang-bang methods, and our numerical results verify that both techniques converge to the same solution.

For the finest grid in Table 5.2, the no bang-bang method with jumps takes about three times more CPU time compared to the same problem with no jumps. This is simply because we need several iterations per timestep to solve the fully implicit discretized equations, including the jump term [17]. Each iteration requires one tridiagonal linear system solve and two FFTs.

The results in Table 5.5 also indicate that the no bang-bang method achieves first-order convergence but not the bang-bang method. To further study this behavior, in Figure 5.4 we show the control curves at $P = 6$ \$/mmBtu, $\tau = \tau^1 = .006$ year as a function of I, obtained using these two methods. We present the control curves produced using one timestep (i.e., $\Delta\tau = \tau^1$) with a coarse space grid, as well as a fine grid solution with many timesteps. From the figure, we can observe that when using finer grids (and more timesteps), both the bang-bang and the no bang-bang methods converge to the same control strategy. In contrast, when a coarse grid and one timestep are used, the control curves produced by both methods differ from the

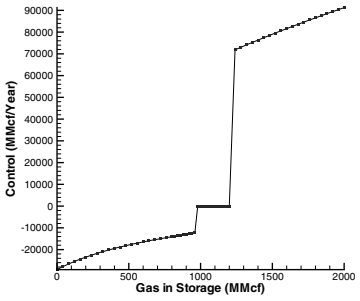

(a) Bang-bang method, $\tau = .006$, one timestep, coarse space grids.

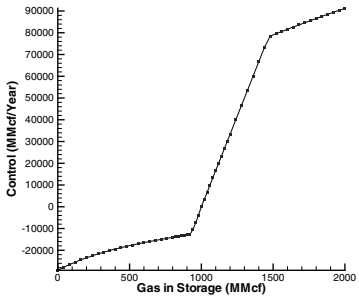

(b) No bang-bang method, $\tau = .006$, one timestep, coarse space grids.

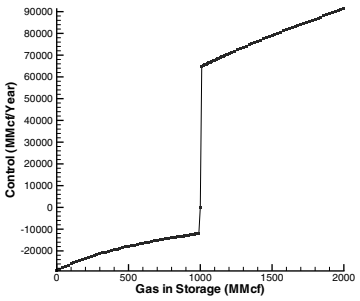

(c) Bang-bang or no bang-bang method, $\tau = .006$, with fine space grid and many timesteps.

FIG. 5.4. *Control curves as a function of gas inventory I obtained at $\tau^1 = .006$ year with gas price $P = 6$ \$/mmBtu. The top panel shows the bang-bang and the no bang-bang methods with one timestep and a coarse space grid. The bottom panel shows the results for both bang-bang and no bang-bang methods, using a fine space grid and many timesteps. Risk-adjusted gas spot price follows the pure mean-reverting process* (5.5). *Payoff function is given in* (5.1). *Input parameters are given in Tables* 5.4 *and* 5.1.

converged controls near $I = 1000$ MMcf (excluding $I = 1000$ MMcf) and hence are not accurate. However, by actually solving the discrete optimization problem in Algorithm 4.9, the no bang-bang method produces a much smoother control curve on coarse grids than does the bang-bang method. Consequently, this would appear to explain why the no bang-bang method is able to generate a smoother solution for the value function as the grid size and timestep size is reduced compared to the bang-bang method.

Figure 5.5 compares the control switching boundary curves obtained before and after incorporating the jump effect when $I = 1000$ MMcf. The figure indicates that the zero control region (the region contained between two boundary curves) resulting

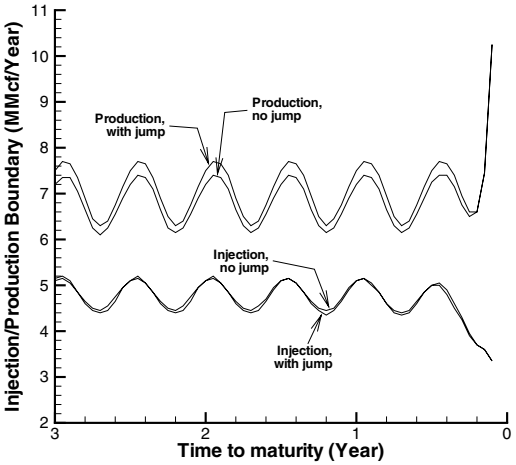

FIG. 5.5. *Control switching boundary curves as a function of time-to-maturity* $\tau = T - t$ *with respect to processes* (5.4) *(without incorporating the jump effect) and* (5.5) *(incorporating the jump effect) at* $I = 1000$ *MMcf, where parameter values for process* (5.5) *are given in Table* 5.4. *Other input parameters are given in Table* 5.1. *Fully implicit timestepping with the no bang-bang method and with constant timesteps is used.*

from the jump diffusion process (5.5) is wider than that resulting from process (5.4). This occurs because, under the jump scenario, the operator is willing to wait for a jump in the gas price and then operate the facility after the jump to obtain more profit, which makes the zero control region wider. In addition, Figure 5.5 shows that the jump effect disappears when the contract is close to maturity because of the fear of revenue loss at maturity due to the payoff structure (5.1).

6. Conclusion. In this paper, we have developed several implicit timestepping schemes based on a semi-Lagrangian method for the valuation of gas storage contracts. This method has the advantage that it avoids the Policy-type iterations required by an implicit finite difference method. Effectively, we replace the Policy iterations by a local optimization problem at each node which involves known values from the previous timestep. This local optimization problem can be solved cheaply. We prove that fully implicit timestepping converges to the viscosity solution of the HJB equation, provided that a strong comparison result holds.

We have used bang-bang and no bang-bang methods to solve the local optimization problems arising from the semi-Lagrangian schemes in order to obtain the optimal operational strategies. The no bang-bang method is a more general method and can be applied to other optimal control problems where the controls are not of bang-bang type.

Our numerical results show that fully implicit and Crank–Nicolson timestepping schemes converge to the same solution. In the case of fully implicit timestepping, at a small extra cost the no bang-bang method generally exhibits smooth first-order convergence for cases where we see nonsmooth convergence for the bang-bang method.

Numerical results also indicate that Crank–Nicolson timestepping can achieve at most first-order convergence. Therefore, we recommend the use of fully implicit timestepping, since convergence to the viscosity solution is guaranteed, assuming that

a strong comparison result holds.

In the future, we plan to price gas storage contracts using more realistic gas spot price models, such as regime switching models. We can also apply our schemes to more complex financial contracts that involve more than one stochastic factor and include complicated revenue structures, such as the valuation of gas powered electrical generation plants [32].

Appendix A. Discrete equation coefficients. In this appendix, we give the coefficients of discrete differential operator $\mathcal{L}_h$ defined in (3.3). For $i = 1$, we impose boundary condition (2.14) as $P \to 0$ by using forward differencing to evaluate the first-order derivative term in (3.1) and setting $V_{PP} = 0$, which results in

$$\alpha_1^n = 0, \qquad \beta_1^n = \frac{\alpha(K(T-\tau^n) - P_1)}{P_2 - P_1}. \tag{A.1}$$

Similarly, for $i = i_{\max}$, the boundary conditions (2.15) as $P \to \infty$ can be imposed by setting

$$\alpha_{i_{\max}}^n = -\frac{\alpha(K(T-\tau^n) - P_{i_{\max}})}{P_{i_{\max}} - P_{i_{\max}-1}}, \qquad \beta_{i_{\max}}^n = 0, \tag{A.2}$$

where $\alpha_{i_{\max}}^n$ and $\beta_{i_{\max}}^n$ are obtained by evaluating the first-order derivative term in (3.1) with backward differencing and setting $V_{PP} = 0$.

Away from $i = 1$ and $i = i_{\max}$, applying the second-order central differencing in the first- and second-order derivative terms in (3.1) leads to the following values of coefficients α_i^n and β_i^n:

$$\begin{aligned} \alpha_{i,central}^n &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_i - P_{i-1})(P_{i+1} - P_{i-1})} - \frac{\alpha(K(T-\tau^n) - P_i)}{P_{i+1} - P_{i-1}}, \\ \beta_{i,central}^n &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_{i+1} - P_i)(P_{i+1} - P_{i-1})} + \frac{\alpha(K(T-\tau^n) - P_i)}{P_{i+1} - P_{i-1}}, \\ & i = 2, \ldots, i_{\max} - 1. \end{aligned} \tag{A.3}$$

If either $\alpha_{i,central}^n$ or $\beta_{i,central}^n$ is negative, the discrete scheme will not be monotone. Monotonicity can be restored by using first-order forward or backward differencing in the first-order derivative term at the problem nodes. Forward differencing produces

$$\begin{aligned} \alpha_{i,forward}^n &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_i - P_{i-1})(P_{i+1} - P_{i-1})}, \\ \beta_{i,forward} &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_{i+1} - P_i)(P_{i+1} - P_{i-1})} + \frac{\alpha(K(T-\tau^n) - P_i)}{P_{i+1} - P_i}, \\ & i = 2, \ldots, i_{\max} - 1, \end{aligned} \tag{A.4}$$

while backward differencing gives

$$\begin{aligned} \alpha_{i,backward}^n &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_i - P_{i-1})(P_{i+1} - P_{i-1})} - \frac{\alpha(K(T-\tau^n) - P_i)}{P_i - P_{i-1}}, \\ \beta_{i,backward}^n &= \frac{(\hat{\sigma}(P_i)P_i)^2}{(P_{i+1} - P_i)(P_{i+1} - P_{i-1})}, \quad i = 2, \ldots, i_{\max} - 1. \end{aligned} \tag{A.5}$$

On the one hand, we want to use central differencing to achieve second-order correctness. On the other hand, we need to maintain α_i^n and β_i^n positive so that the scheme is monotone. Consequently, we decide on central or forward/backward discretization at each node (P_i, I_j), $i = 2, \ldots, i_{\max} - 1$, $j = 1, \ldots, j_{\max}$, based on the strategy described in [17]. The strategy guarantees that both α_i^n and β_i^n are nonnegative. Equations (A.3) imply that $\alpha_{i,central}^n$ and $\beta_{i,central}^n$ are positive if P_i is close to $K(T-\tau^n)$, which is the equilibrium price of the risk-adjusted natural gas spot price at $t = T - \tau^n$. Hence the strategy in [17] will use central differencing for those nodes P_i close to the equilibrium price. These nodes are in the region of interest, since the gas spot price will not stray too far away from the mean-reversion equilibrium price. Consequently, the use of a low-order differencing scheme for nodes far away from the equilibrium price should not result in poor convergence for the nodes near the equilibrium price.

Appendix B. Sketch of proof for Lemma 4.5. The proof for Lemma 4.5 mainly depends on the following two lemmas.

LEMMA B.1. *Let $C_j^{n+1} \subseteq C\big(\mathcal{I}_{j(i,n+1)}^n\big) \times C(I_j)$ be the set of admissible controls such that $\mathcal{I}_{j(i,n+1)}^n$ calculated from*

$$\mathcal{I}_{j(i,n+1)}^n = I_j - (1-\theta)\Delta\tau^n\big(\zeta_{i,j}^{n+1} + a(\zeta_{i,j}^{n+1})\big) - \theta\Delta\tau^n\big(\zeta_{i,j}^{n} + a(\zeta_{i,j}^{n})\big) \tag{B.1}$$

is bounded within $[0, I_{\max}]$, where $C\big(\mathcal{I}_{j(i,n+1)}^n\big), C(I_j)$ are defined based on (2.7). *Then by taking $\Delta\tau^n$ sufficiently small, we have*

$$C_j^{n+1} = C\big(\mathcal{I}_{j(i,n+1)}^n\big) \times C(I_j), \tag{B.2}$$

provided that $c_{\max}$, $c_{\min}$, and $a(c)$ satisfy (2.1), (2.3), *and* (2.5), *respectively.*

Proof. The proof is outlined in the following two steps. See [13, Appendix D] for a more detailed proof. In the first step, we consider the case where $\epsilon^* \le I_j \le I_{\max} - \epsilon^*$, where $\epsilon^* \ll I_{\max}$ is independent of $h_{\min}$. In other words, we exclude a small strip of finite size near $I = 0$ and $I = I_{\max}$. Using the fact that any control values are bounded, we can show that condition (B.2) holds if $\Delta\tau^n < D$, where D is a $O(1)$ constant.

In the second step, we consider the nodes in the boundary strips $I_j < \epsilon^*$ and $I_j > I_{\max} - \epsilon^*$. Conducting a similar analysis as in the first step seems to require an additional condition on $\Delta\tau^n/h_{\min}$. However, using (2.1), (2.3), and (2.5), it can be shown that condition (B.2) follows if $\Delta\tau^n \le \text{const.}\sqrt{h_{\min}}$. However, this is a weaker condition than the assumption (3.22). Therefore, for both cases, condition (B.2) follows for $\Delta\tau^n$ sufficiently small. ∎

LEMMA B.2. *Suppose that $F(c_1, c_2)$ and $H(c_1, c_2)$ are bounded functions defined in some domain $c_1 \in C_1$, $c_2 \in C_2$. Then there exists a bounded function $Q(h_{\max})$, where $Q(h_{\max}) \le \max_{c_1,c_2} |H(c_1, c_2)|$, such that*

$$\min_{c_1,c_2}\{F(c_1,c_2) + H(c_1,c_2)h_{\max}\} = \min_{c_1,c_2}\{F(c_1,c_2)\} + h_{\max}Q(h_{\max}). \tag{B.3}$$

Proof. According to [23], we have that

$$\begin{aligned}
&\min_{c_1,c_2}\{H(c_1,c_2)h_{\max}\} \le \min_{c_1,c_2}\{F(c_1,c_2) + H(c_1,c_2)h_{\max}\} - \min_{c_1,c_2}\{F(c_1,c_2)\},\\
&\max_{c_1,c_2}\{H(c_1,c_2)h_{\max})\} \ge \min_{c_1,c_2}\{F(c_1,c_2) + H(c_1,c_2)h_{\max}\} - \min_{c_1,c_2}\{F(c_1,c_2)\}.
\end{aligned} \tag{B.4}$$

Equation (B.3) directly follows from the above inequalities. ∎

After presenting Lemmas B.1 and B.2, we sketch the proof of Lemma 4.5 as follows. We define vector ϕ^n with components $[\phi^n]_{i,j} = \phi(P_i, I_j, \tau^n)$. We first consider the discrete function $G^{n+1}_{i,j}$ in (3.23). For convenience, we define

$$W\big(P_i, I_j, \tau^{n+1}; c\big) = \big(\phi_\tau + (c + a(c))\phi_I - (c - a(c))P - \mathcal{L}\phi\big)^{n+1}_{i,j}. \tag{B.5}$$

Applying Taylor series expansion to $G^{n+1}_{i,j}$ and writing the result in terms of notation (B.5) gives

$$\begin{aligned}
&G^{n+1}_{i,j}\big(h_{\max}, \phi^{n+1}_{i,j}, \{\phi^{n+1}_{k,j}\}_{k\neq i}, \{\phi^n_{i,j}\}\big) \\
&= \min_{(\zeta^n_{i,j}, \zeta^{n+1}_{i,j})\in C^{n+1}_j} \Big\{(1-\theta)W\big(P_i, I_j, \tau^{n+1}; \zeta^{n+1}_{i,j}\big) + \theta W\big(P_i, I_j, \tau^{n+1}; \zeta^n_{i,j}\big) \\
&\qquad\qquad + O\left(\Delta\tau_{\max} + \Delta P_{\max} + (\Delta I_{\max})^2/\Delta\tau_{\min} + (\Delta I_{\max})^2\right)\Big\} \\
&= \min_{(\zeta^n_{i,j}, \zeta^{n+1}_{i,j})\in C^{n+1}_j} \big\{(1-\theta)W\big(P_i, I_j, \tau^{n+1}; \zeta^{n+1}_{i,j}\big) + \theta W\big(P_i, I_j, \tau^{n+1}; \zeta^n_{i,j}\big) \\
&\qquad\qquad + O(h_{\max})\big\},
\end{aligned} \tag{B.6}$$

where we assume that the interpolation error in (3.19), due to operation Φ^{n+1}, is $O((\Delta I_{\max})^2)$—i.e., linear or higher order interpolation is used—and we use (3.22) to unify the mesh size/timestep size using $h_{\max}$.

According to Lemma B.1, by taking $\Delta\tau^n$ sufficiently small, given (2.1), (2.3), and (2.5), we can relax the constraint $(\zeta^n_{i,j}, \zeta^{n+1}_{i,j}) \in C^{n+1}_j$ in above equation (B.6) to $\zeta^{n+1}_{i,j} \in C(I_j)$, $\zeta^n_{i,j} \in C\big(\mathcal{I}^n_{j(i,n+1)}\big)$.

The above argument together with Lemma B.2 allows us to write (B.6) as

$$\begin{aligned}
&G^{n+1}_{i,j}\big(h_{\max}, \phi^{n+1}_{i,j}, \{\phi^{n+1}_{k,j}\}_{k\neq i}, \{\phi^n_{i,j}\}\big) \\
&= (1-\theta) \min_{\zeta^{n+1}_{i,j}\in C(I_j)} \big\{W\big(P_i, I_j, \tau^{n+1}; \zeta^{n+1}_{i,j}\big)\big\} \\
&\quad + \theta \min_{\zeta^n_{i,j}\in C(\mathcal{I}^n_{j(i,n+1)})} \big\{W\big(P_i, I_j, \tau^{n+1}; \zeta^n_{i,j}\big)\big\} + Q(h_{\max})h_{\max},
\end{aligned} \tag{B.7}$$

where Q is a bounded function. Note that as $h_{\max} \to 0$, then $\Delta\tau^n \to 0$, and hence $C\big(\mathcal{I}^n_{j(i,n+1)}\big) \to C(I_j)$ in (B.7). Now substitute (B.7) into the left-hand side of (4.3), and it is straightforward to see that the equality in (4.3) follows as $h_{\max} \to 0$.

REFERENCES

[1] H. Ahn, A. Danilova, and G. Swindle, *Storing Arb*, Wilmott Magazine, 1 (2002), pp. 78–83.

[2] G. Barles, *Convergence of numerical schemes for degenerate parabolic equations arising in finance*, in Numerical Methods in Finance, L. C. G. Rogers and D. Talay, eds., Cambridge University Press, Cambridge, UK, 1997, pp. 1–21.

[3] G. Barles, *Nonlinear Neumann boundary conditions for quasilinear degenerate elliptic equations and applications*, J. Differential Equations, 154 (1999), pp. 191–224.

[4] G. Barles and J. Burdeau, *The Dirichlet problem for semilinear second-order degenerate elliptic equations and applications to stochastic exit time control problems*, Comm. Partial Differential Equations, 20 (1995), pp. 129–178.

[5] G. Barles and E. Jakobsen, *Error bounds for monotone approximation schemes for parabolic Hamilton–Jacobi–Bellman equations*, Math. Comp., 27 (2007), pp. 1861–1893.
[6] G. Barles and E. Rouy, *A strong comparison result for the Bellman equation arising in stochastic exit time control problems and its applications*, Comm. Partial Differential Equations, 23 (1998), pp. 1945–2033.
[7] G. Barles and P. E. Souganidis, *Convergence of approximation schemes for fully nonlinear equations*, Asymptot. Anal., 4 (1991), pp. 271–283.
[8] C. Barrera-Esteve, F. Bergeret, C. Dossal, E. Gobet, A. Meziou, R. Munos, and D. Reboul-Salze, *Numerical methods for the pricing of swing options: A stochastic control approach*, Methodol. Comput. Appl. Probab., 8 (2006), pp. 517–540.
[9] R. Bermejo, *Analysis of a class of quasi-monotone and conservative semi-Lagrangian advection schemes*, Numer. Math., 87 (2001), pp. 597–623.
[10] M. Bourgoing, *Viscosity Solutions of Fully Nonlinear Second Order Parabolic Equations with l^1 Dependence in Time and Neumann Boundary Conditions*, working paper, Université de Tours, Tours, France, 2004.
[11] R. Carmona and M. Ludkovski, *Gas storage and supply guarantees: An optimal switching approach*, Management Sci., submitted.
[12] S. Chaumont, *A Strong Comparison Result for Viscosity Solutions to Hamilton–Jacobi–Bellman Equations with Dirichlet Condition on a Non-smooth Boundary and Application to Parabolic Problems*, Preprint submitted to Elsevier Science, Université Henri Poincaré I, 2004.
[13] Z. Chen and P. A. Forsyth, *A Semi-Lagrangian Approach for Natural Gas Storage Valuation and Optimal Operation*, Technical report, David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, 2006.
[14] M. G. Crandall, H. Ishii, and P. L. Lions, *User's guide to viscosity solutions of second order partial differential equations*, Bull. Amer. Math. Soc., 27 (1992), pp. 1–67.
[15] Y. d'Halluin, P. Forsyth, and K. Vetzal, *Robust numerical methods for contingent claims under jump diffusion processes*, IMA J. Numer. Anal., 25 (2005), pp. 65–92.
[16] Y. d'Halluin, P. A. Forsyth, and G. Labahn, *A penalty method for American options with jump diffusion processes*, Numer. Math., 97 (2004), pp. 321–352.
[17] Y. d'Halluin, P. A. Forsyth, and G. Labahn, *A semi-Lagrangian approach for American Asian options under jump diffusion*, SIAM J. Sci. Comput., 27 (2005), pp. 315–345.
[18] A. Dixit and R. Pindyck, *Investment Under Uncertainty*, Princeton University Press, Princeton, NJ, 1994.
[19] J. Douglas, Jr., and T. F. Russell, *Numerical methods for convection-dominated diffusion problems based on combining the method of characteristics with finite element or finite difference procedures*, SIAM J. Numer. Anal., 19 (1982), pp. 871–885.
[20] M. Falcone and R. Ferretti, *Discrete time high-order schemes for viscosity solutions of Hamilton-Jacobi-Bellman equations*, Numer. Math., 67 (1994), pp. 315–344.
[21] M. Falcone and R. Ferretti, *Convergence analysis for a class of high-order semi-Lagrangian advection schemes*, SIAM J. Numer. Anal., 35 (1998), pp. 909–940.
[22] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, Springer, New York, 2006.
[23] P. A. Forsyth and G. Labahn, *Numerical methods for controlled Hamilton–Jacobi–Bellman PDEs in finance*, J. Comp. Fin., to appear.
[24] E. R. Jakobsen, *On the rate of convergence of approximation schemes for Bellman equations associated with optimal stopping time problems*, Math. Models Methods Appl. Sci., 13 (2003), pp. 613–644.
[25] H. Li, *Adaptive Wavelet Collocation Methods for Option Pricing PDEs*, Ph.D. Thesis, Department of Mathematics, University of Calgary, Calgary, AB, 2006.
[26] H. Pham, *On some recent aspects of stochastic control and their applications*, Probability Surveys, 2 (2005), pp. 506–549.
[27] D. Pilopović, *Energy Risk*, McGraw–Hill, New York, 1998.
[28] O. Pironneau, *On the transport diffusion algorithm and its applications to the Navier-Stokes equations*, Numer. Math., 38 (1982), pp. 309–332.
[29] D. M. Pooley, P. A. Forsyth, and K. R. Vetzal, *Numerical convergence properties of option pricing PDEs with uncertain volatility*, IMA J. Numer. Anal., 23 (2003), pp. 241–267.
[30] R. Rannacher, *Finite element solution of diffusion problems with irregular data*, Numer. Math., 43 (1984), pp. 309–327.
[31] M. Thompson, M. Davison, and H. Rasmussen, *Natural Gas Storage Valuation and Optimization: A Real Options Application*, working paper, University of Western Ontario, London, ON, 2003.

[32] M. Thompson, M. Davison, and H. Rasmussen, *Valuation and optimal operation of electric power plants in competitive markets*, Oper. Res., 52 (2004), pp. 546–562.
[33] L. Triegeorgis, *Real Options: Managerial Flexibility and Strategy for Resource Allocation*, MIT Press, Boston, 1996.
[34] T. Ware, *Swing options in a mean-reverting world*, presentation at Stochastic Calculus and Its Applications to Quantitative Finance and Electrical Engineering, a conference in honor of Robert Elliott, Calgary, July 2005.
[35] T. Ware and H. Li, *Swing options with continuous exercise*, presentation at the Canadian Mathematical Society Meeting, Calgary, June 2006.
[36] R. Zvan, P. A. Forsyth, and K. R. Vetzal, *Discrete Asian barrier options*, J. Comput. Finance, 3 (1999), pp. 41–68.