



Volatility uncertainty quantification in a stochastic control problem applied to energy

Francisco Bernal, Emmanuel Gobet, Jacques Printems

► To cite this version:

Francisco Bernal, Emmanuel Gobet, Jacques Printems. Volatility uncertainty quantification in a stochastic control problem applied to energy. *Methodology and Computing in Applied Probability*, 2019, 22 (1), pp.135-159. 10.1007/s11009-019-09692-x . hal-01784095

HAL Id: hal-01784095

<https://hal.science/hal-01784095>

Submitted on 3 May 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Volatility uncertainty quantification in a stochastic control problem applied to energy*

Francisco Bernal[†] Emmanuel Gobet[‡] Jacques Printems[§]

February 25, 2018

Abstract

This work designs a methodology to quantify the uncertainty of a volatility parameter in a stochastic control problem arising in energy management. The difficulty lies in the non-linearity of the underlying scalar Hamilton-Jacobi-Bellman equation. We proceed by decomposing the unknown solution on a Hermite polynomial basis (of the unknown volatility), whose different coefficients are solution to a system of non-linear PDEs of the same kind. Numerical tests show that computing the first basis elements may be enough to get an accurate approximation with respect to the uncertain volatility parameter. We experiment the methodology in the context of swing contract (energy contract with flexibility in purchasing energy power), this allows to introduce the concept of Uncertainty Value Adjustment (UVA), whose aim is to value the risk of misspecification of the volatility model.

Key words: chaos expansion, uncertainty quantification, stochastic control, stochastic programming, Swing options, Monte Carlo simulations.

MSC 2010: 93Exx, 62L20, 41A10, 90C15, 49L20.

*This research is part of the Chair *Financial Risks* of the *Risk Foundation*, the *Finance for Energy Market Research Centre* (FiME) and the ANR project *CAESARS* (ANR-15-CE05-0024).

[†]CMAP, Ecole Polytechnique, email: francisco.bernal@polytechnique.edu

[‡]CMAP, Ecole Polytechnique, email: emmanuel.gobet@polytechnique.edu

[§]LAMA, Université Paris-Est Créteil, email: printems@u-pec.fr

1 Introduction

1.1 Statement of the problem

Swing or Take-Or-Pay contracts are very common in energy markets. When the underlying commodity is gas, they give the possibility to the buyer of such a contract to buy/exchange the gas on the spot market (or the forward market) at a given price, with a flexibility on the cumulative volume on a given period, say $[0, T]$. The valuation problem can be split into two steps: identifying the model from data, and computing the swing option price given the model. The second step is usually quite demanding since it requires to solve a non-linear Hamilton-Jacobi-Bellman (HJB) equation: for some related numerical methods, see [CSK01, JRT04, Kep04, BEBD⁺06, BBP09]. Here we address the problem of swing valuation when the volatility of the model is imperfectly identified.

As a model, we consider the Schwartz-Smith model [SS00] simplified to a single factor (see also [Bre91]-[Sch97]): namely, we assume that the forward gas price $F(s, t)$ (fixed at time s for receiving and paying one unit of gas at time $t > s$) obeys the one-factor dynamics (for any maturity t)

$$\frac{dF(s, t)}{F(s, t)} = \sigma e^{-\alpha(t-s)} dW_s, \quad 0 \leq s < t, \quad (1.1)$$

given a initial forward curve ($F(0, t) : t \geq 0$). The parameter $\sigma > 0$ is the volatility and $\alpha > 0$ is the mean-reversion coefficient that describes the rate at which the short-term deviations are expected to disappear (see (2.2) later, for the link with mean-reverting Ornstein-Uhlenbeck process).

The statistical inference of (α, σ) is usually handled with Maximum Likelihood estimation or Kalman filtering [Sch97, SS00, MT02]; standard statistical results [Har89, Rao99] show that with a large number of data, the estimators are asymptotically normal (for the Local Asymptotic Normality property, see [Gob02]). Therefore, naturally, one may consider that the estimation errors on (α, σ) are approximately Gaussian, with small covariance/variances. Here to simplify the analysis, we consider that α is known and only σ is subjected to estimation error. In other words, we have

$$\sigma \stackrel{d}{\approx} \mathcal{N}(\mu, \nu^2), \quad (1.2)$$

with small variance $\nu > 0$ and some $\mu > 0$. This comes from a Central Limit Theorem result where typically ν^2 is proportional to the inverse of the number of available data.

The model-certain price $v^\sigma(0)$ at time 0 of the swing contract depends naturally on σ , and we wish now to account for the variability of price with respect to σ ; in other words, to quantify the uncertainty in σ . Typically our aim is to efficiently estimate the distribution of the random variable $v^\sigma(0)$. Because $v^\sigma(0)$ solves a non-linear equation (for each given σ), it raises new issues that have not yet been addressed before.

From the point of view of financial economics, evaluating the mean excess between the uncertain swing option price $v^\sigma(\sigma)$ and the deterministic price $v^\sigma(0)$ is helpful to account for the risk of model uncertainty. This is similar to Credit Value Adjustment in finance, and other related adjustments, collectively called xVA (see [Bas15]), corresponding (in the case of CVA) to accounting for the possibility of a counter-party's default within a bilateral contract. In our case, $\mathbb{E}[(v^\sigma(\sigma) - z)_+]$ with $z \geq v^\sigma(0)$ is referred to as UVA (Uncertainty Value Adjustment). The level z could be taken as $v^\sigma(0) = \mathbb{E}[v^\sigma(\sigma)]$, or as a given quantile of $v^\sigma(\sigma)$. This measure may serve to compute a deposit in order to face a quite unfavorable pricing of the swing option. Some illustrations are presented in Section 6.

1.2 Background results in uncertainty quantification

Many numerical strategies for uncertainty quantification are possible, of which we single out the most usual ones. For definiteness, we assume that a set of κ uncertain parameters σ is being propagated by a PDE, which are usually modeled as independent random distributions, and we say the random space is κ -dimensional. In the case of an uncertain function, modeled by a random field, dimension reduction is typically performed first via Karhunen-Loève expansions (KLE) truncated to a finite number κ of random parameters. Here are the different numerical strategies.

- Simulation-based methods, consisting in sampling M different σ according to (1.2) and for each of them, solving $(v^{\sigma^{(m)}}(0) : 1 \leq m \leq M)$. Sample statistics can then be performed on the ensemble of simulations. These methods are straightforward to implement provided that a deterministic solver for the PDE with fixed parameters is available (hence they are called non-intrusive). Moreover, they can handle arbitrary uncertainty distributions and high-dimensional uncertainty spaces with no particular modification. On the other hand, the convergence rate is very slow (typically as $M^{-1/2}$), a problem which is compounded if each deterministic sample is costly to compute. Strategies to reduce the statistical error (and thus to speed up simulations) include Latin hypercube [Loh96], quasi Monte Carlo [Nie92] and the Multilevel method [TSGU13].

- Local expansion-based methods, consisting in deriving Taylor/perturbation methods, taking advantage (or assuming) that the uncertain parameter σ fluctuates little. See [KH92]. It boils down to a sensitivity analysis which appears to be quite delicate in the case of HJB equation, because of the non-linear property of the equation (see (2.4) later).
- Generalized Polynomial Chaos (gPC) methods [Xiu09][LK10], which involve a functional expansion (the gPC expansion) of the solution in a basis in the random space—ideally orthogonal with respect to the joint distribution of parameters. There are two main approaches to gPC, differing in how the residual to the PDE is treated. In the first one—stochastic collocation—the residual of the gPC expansion is forced to vanish on a predefined set of nodes in the random space. This involves the discretization of the random space, which can be costly in high dimensions ($\kappa \gtrsim 4$)—although sparse grids can be sometimes used—but has the advantage of being non-intrusive. The second approach is called stochastic Galerkin, where the residual is forced to be orthogonal to each of the basis functions in the random space retained in the gPC expansion. This method is usually considered more accurate than stochastic collocation and does not involve grids in the random space; but the projection of the residual results in a system of κ PDEs which are different from the original problem (hence it is an intrusive method).

In this work we follow the last approach using polynomial chaos expansion (PCE) based on Gaussian noise (due to (1.2)). To the best of our knowledge, the above techniques have been investigated in the case of linear problems (like [MR98]) and scarcely in the case of nonlinear PDEs (in fluid mechanics see [MR05], in stochastic optimization and HJB equations see [BL14, HS14], both cases are not related to Uncertainty Quantification strictly speaking).

2 Model and polynomial chaos expansion

2.1 Modeling the energy contract: HJB equation with uncertainty

To better represent (1.1), define

$$X_s = \sigma e^{-\alpha s} \int_0^s e^{\alpha u} dW_u, \quad (2.1)$$

which is an Ornstein-Uhlenbeck process, starting from 0, and with zero long-term average:

$$dX_s = -\alpha X_s ds + \sigma dW_s, \quad X_0 = 0.$$

The process X is Gaussian and its limiting distribution as time goes to infinity is $\mathcal{N}(0, \frac{\sigma^2}{2\alpha})$. Then, by the Itô formula we have

$$F(s, t) = F(0, t) \exp \left[e^{-\alpha(t-s)} X_s - \frac{\sigma^2}{4\alpha} \left(e^{-2\alpha(t-s)} - e^{-2\alpha t} \right) \right]. \quad (2.2)$$

The gas spot price is

$$S_t = F(t, t) = F(0, t) \exp \left[X_t - \frac{\sigma^2}{4\alpha} \left(1 - e^{-2\alpha t} \right) \right]. \quad (2.3)$$

The swing contract has an expiration equal $T > 0$, and the fixed price to buy gas is given by the strike $K > 0$. At any date, the swing owner has the right to buy q_t units of gas at the price K , so that his instantaneous profit is

$$f(t, X_t, q_t) = q_t(S_t - K).$$

Here the volume flexibility is allowed in the range $q_t \in [q_m, q_M] := Q \subset [0, +\infty)$ at any time $t \in [0, T]$. The cumulative gas withdrawn is

$$Y_t = \int_0^t q_s ds.$$

The value function is defined as

$$v^\sigma(t, x, y) = \sup_{(q_s)_{t \leq s \leq T}, q_s \in Q} \mathbb{E} \left[P(X_T, Y_T) + \int_t^T f(s, X_s, q_s) ds \mid X_t = x, Y_t = y \right] \quad (2.4)$$

where the expectation is taken under the risk-neutral measure defining the forward contract/spot price. The term P plays the role of a penalty forcing the cumulative withdraw Y_T to finish in a given range (depending on the swing contract), typically we can take P with polynomial growth in its variables. Under smoothness assumptions on v , we know that v solves the following Hamilton-Jacobi-Bellman PDE:

$$\begin{cases} \partial_t v^\sigma + \sup_{q \in Q} (\mathcal{L}^{\sigma, q} v^\sigma + f) = 0, & (t, x, y) \in [0, T] \times \mathbb{R} \times \mathbb{R}, \\ v^\sigma(T, x, y) = P(x, y), \end{cases}$$

where

$$\mathcal{L}^{\sigma,q}\Psi(t, x, y) = \frac{\sigma^2}{2}\partial_{xx}^2\Psi(t, x, y) - \alpha x\partial_x\Psi(t, x, y) + q\partial_y\Psi(t, x, y).$$

By changing $t \rightarrow T-t$, we get another PDE ($v^\sigma(t, \cdot) = u^\sigma(T-t, \cdot)$) with initial condition:

$$\begin{cases} \partial_t u^\sigma = \frac{\sigma^2}{2}\partial_{xx}^2 u^\sigma - \alpha x\partial_x u^\sigma + \sup_{q \in Q} q \left(\partial_y u^\sigma + s(t, \sigma)e^x - K \right), \\ u^\sigma(0, x, y) = P(x, y), \end{cases} \quad (2.5)$$

where $s(t, \sigma) = F(0, T-t)e^{-\frac{\sigma^2}{4\alpha}(1-e^{-2\alpha(T-t)})}$ (see Equation (2.3)). This is a degenerate diffusion along y compounded by advection in that direction (in the sup term).

Boundary conditions. Equation (2.5) is defined on the halfplane $y \geq 0$. However, the solution is only needed in a bounded region I around the origin where contract scenarios with significant probability occur. For computational purposes, a finite computational domain $D = [x_m, x_M] \times [0, y_M]$ with $x_m < 0 < x_M$ must be defined which contains I . The boundary conditions (BCs) on ∂D are to be chosen in such a way that they do not affect the solution inside I ; and they are compatible with the initial condition P .

Along the y direction, (2.5) is purely hyperbolic and, since $q \geq q_m \geq 0$, information travels downward. Consequently, no BC is required at the downwind side of the boundary ($y = 0$), while the BC on the upwind side cannot affect I as long as it is farther away than $q_M T$, i.e. if $y_M > \sup_{y \in I} y + q_M T$.

Along the x direction, (2.5) is diffusive and information travels instantaneously; the goal is to place the BC far enough so that its effect inside I is negligible. In order to get a rough estimate of $x_L := \max(-x_m, x_M)$, let us heuristically split (2.5) into three parts: a diffusion term $\frac{\sigma^2}{2}\partial_{xx}^2 u^\sigma$; an advection term $-\alpha x\partial_x u^\sigma$, which advects information away from I ; and a nonlinear term, which we neglect. Based on the fundamental solution of the diffusion part only, the effect $b > 0$ of a Dirichlet BC $u(t', x_L, y)$ on a point $(x_I, y_I) \in I$ at time $t > t'$ ("boundary pollution") is assumed bounded as

$$b \lesssim \frac{|u(t', x_L, y)|}{\sqrt{2\pi\sigma^2(t-t')}} \exp\left(-\frac{|x_I - x_L|^2}{2\sigma^2(t-t')}\right). \quad (2.6)$$

In this way, given a prescribed tolerance b , a rough estimate of x_L can be obtained.

Let us define

$$Z^\sigma(t, x, y) := \partial_y u^\sigma(t, x, y) + s(t, \sigma)e^x - K, \quad (2.7)$$

which may be denoted simply by Z^σ when there is no ambiguity. Then

$$\sup_{q \in Q} qZ = q_M(Z)_+ - q_m(Z)_- = Z(q_M \mathbf{1}_{Z \geq 0} + q_m \mathbf{1}_{Z < 0})$$

and the HJB equation is

$$\begin{cases} \partial_t u^\sigma = \frac{\sigma^2}{2} \partial_{xx}^2 u^\sigma - \alpha x \partial_x u^\sigma + Z^\sigma (q_M \mathbf{1}_{Z^\sigma \geq 0} + q_m \mathbf{1}_{Z^\sigma < 0}), \\ u^\sigma(0, x, y) = g(y). \end{cases} \quad (2.8)$$

2.2 Hermite Polynomial chaos expansion

From (1.2), we assume an equality and we may write

$$\sigma^2(\xi) = (\mu + \nu \xi)^2, \quad (\mathbf{Q})$$

leading to the approximation (1.2) as $\nu \rightarrow 0$. Here ξ is a standard Gaussian random variable $\mathcal{N}(0, 1)$, with density

$$p(\xi) = \frac{1}{\sqrt{2\pi}} e^{-\xi^2/2}.$$

The uncertain parameter being related to Gaussian distribution, we derive a PCE using Hermite polynomials. This is the purpose of the subsequent presentation.

2.2.1 Hermite polynomials

Consider the orthogonal basis of Hermite polynomials [LK10, Appendix B.1.2 p.502]: they are given by the Rodrigues formula

$$\phi_n(\xi) = (-1)^n e^{\frac{\xi^2}{2}} \partial_\xi^n [e^{-\xi^2/2}], \quad n \geq 0.$$

Therefore,

$$\begin{cases} \phi_0(\xi) = 1, & \phi_1(\xi) = \xi, & \phi_2(\xi) = \xi^2 - 1, \\ \phi_3(\xi) = \xi^3 - 3\xi, & \phi_4(\xi) = \xi^4 - 6\xi^2 + 3, \\ \phi_5(\xi) = \xi^5 - 10\xi^3 + 15\xi, & \phi_6(\xi) = \xi^6 - 15\xi^4 + 45\xi^2 - 15, \end{cases}$$

and

$$\mathbb{E}[\phi_n(\xi) \phi_m(\xi)] = \int_{-\infty}^{+\infty} \phi_n(\xi) \phi_m(\xi) p(\xi) d\xi = n! \delta_{nm}. \quad (2.9)$$

In particular, we set

$$V_k := \mathbb{E} [\phi_k^2(\xi)] = k!. \quad (2.10)$$

These polynomials satisfy a three-term relation, as any family of orthogonal polynomials; here it takes the form:

$$\begin{cases} \xi \phi_n(\xi) = \phi_{n+1}(\xi) + n\phi_{n-1}(\xi), & n \geq 1, \\ \phi_1(\xi) = \xi\phi_0(\xi). \end{cases} \quad (2.11)$$

From this, we deduce

$$\begin{cases} \xi^2 \phi_n(\xi) = \phi_{n+2}(\xi) + (2n+1)\phi_n(\xi) + n(n-1)\phi_{n-2}(\xi), & n \geq 2, \\ \phi_3(\xi) = (\xi^2 - 3)\phi_1(\xi), \end{cases} \quad (2.12)$$

that we will use subsequently to derive the system of non-linear PDEs to solve.

2.2.2 Chaos expansion, Monte-Carlo evaluation of the distribution of $u^\sigma(\cdot)$

Since $u^{\sigma(\xi)} : \xi \in \mathbb{R} \mapsto L_2(L_2([0, T] \times D_L; \mathbb{R}), p(\xi)d\xi)$, it can be decomposed onto the Hermite basis with coefficients in $L_2([0, T] \times D_L; \mathbb{R})$ (see [LK10]):

$$u^{\sigma(\xi)}(\cdot) = \sum_{k=0}^{+\infty} \tilde{u}_k(\cdot) \phi_k(\xi), \quad (2.13)$$

with (for any t, x, y)

$$\begin{cases} \tilde{u}_k(t, x, y) = \mathbb{E} [u^{\mu+\nu\xi}(t, x, y) \phi_k(\xi)] / k!, \\ \mathbb{E} [|u^{\sigma(\xi)}(t, x, y)|^2] = \sum_{k=0}^{+\infty} \tilde{u}_k^2(t, x, y) k!. \end{cases} \quad (2.14)$$

Thus, in the decomposition (2.13), $(\tilde{u}_k^2(t, x, y) k!)_{k \geq 0}$ must form a convergent series for almost every $(t, x, y) \in [0, T] \times D_L$. Following this PCE approach, our goal is now to identify the sequence of function coefficients $\tilde{u}_k(\cdot), k \geq 0$.

Once the coefficients identified (see Subsection 2.2.3 below), note that it is very easy (and efficient) to sample the distribution of a functional of $u^\sigma(\cdot)$. For instance, if one is interested by the distribution of a single value $u^\sigma(t, x, y)$, it is enough to sample i.i.d. standard Gaussian ξ_1, \dots, ξ_M and to evaluate the approximation $\sum_{k=0}^n \tilde{u}_k(t, x, y) \phi_k(\xi)$ as those sample points $\{\xi = \xi_m, m = 1, \dots, M\}$ and for some large order n : it will provide an empirical measure of $u^\sigma(t, x, y)$ under the uncertainty measure for σ . If

instead of a single value one needs a functional of the solution as a function of $(t, x, y) \in [0, T] \times D_L$, the methodology is unchanged. This approach is significantly faster than the naive one consisting to solve the whole PDE for each sampled $\xi = \xi_m$, since only the $\phi_k(\xi)$ have to be computed and the $\{\tilde{u}_k : 0 \leq k \leq n\}$ do not depend anymore on the noise.

If only the expectation of $u^\sigma(t, x, y)$ is required, it is enough to consider $\tilde{u}_0(t, x, y)$ in view of (2.14), so that no extra simulation is needed. The variance is also obtained simply, without extra effort, by using the relation

$$\mathbb{V}\text{ar}(u^{\sigma(\xi)}(t, x, y)) = \sum_{k=1}^{+\infty} \tilde{u}_k^2(t, x, y) k! \approx \sum_{k=1}^n \tilde{u}_k^2(t, x, y) k!.$$

Higher order moments can be computable explicitly too, see [LK10, Appendix C]. This makes this PCE approach potentially much cheaper than the crude Monte-Carlo method.

2.2.3 Representation of the chaos expansion from the HJB equation

We truncate (Galerkin projection) the expansion (2.13) at the truncation level $n \geq 0$:

$$u^{\sigma(\xi)}(t, x, y) \approx u_n^\xi(t, x, y) := \sum_{k=0}^n \tilde{u}_k(t, x, y) \phi_k(\xi).$$

The above residual at order n is orthogonal to $\{\phi_i(\xi)\}_{i=0}^n$. From (2.8) and (2.14), we obtain a PC system ($k = 0, 1, \dots, n$)

$$\mathbb{E} [\phi_k(\xi) \partial_t u_n^\xi(t, x, y)] = \mathbb{E} \left[\phi_k(\xi) \frac{\sigma^2(\xi)}{2} \partial_{xx}^2 u_n^\xi(t, x, y) \right] - \mathbb{E} [\phi_k(\xi) \alpha x \partial_x u_n^\xi(t, x, y)] \quad (2.15)$$

$$+ \mathbb{E} \left[\phi_k(\xi) Z_n^\xi(t, x, y) \left(q_M \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} + q_m \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right) \right], \quad (2.16)$$

where we set (similarly to (2.7))

$$Z_n^\xi(t, x, y) = Z(t, x, y, u_n^\xi(t, x, y), \xi) = \partial_y u_n^\xi(t, x, y) + s(t, \sigma(\xi)) e^x - K.$$

We may add a similar treatment of the initial and boundary conditions. The following transformations of (2.15)-(2.16) are obtained using extensively the orthogonality relation (2.9) and the definition of u_n , assuming all the necessary smoothness on u (and thus on \tilde{u}_k).

▷ **Term with time-derivative.** The time-derivative and the first-order derivative in y (available in y) can be transformed into semi-Lagrangian derivative. In order to get rid of advection along a degenerate direction of the diffusion (y), let us define the semi-Lagrangian derivative along a characteristic $\Gamma(t) : y(t) = y(0) + qt, t \geq 0$, as

$$D_t^{(q)}u(t, x, y) := D_t^{(q)}u(t, x, y \in \Gamma(t)) = \partial_t u(t, x, y) + q\partial_y u(t, x, y).$$

It can be discretized for instance as

$$D_t^{(q)}u(t, x, y) \approx \frac{u(t, x, y) - u(t - \Delta t, x, y - q\Delta t)}{\Delta t}.$$

Set

$$\begin{aligned} W_k^+(t, x, y) &= \mathbb{E} \left[\phi_k(\xi) [s(t, \sigma(\xi))e^x - K] \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right], \\ W_k^-(t, x, y) &= \mathbb{E} \left[\phi_k(\xi) [s(t, \sigma(\xi))e^x - K] \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right], \\ H_{jk}^+(t, x, y) &= \mathbb{E} \left[\phi_j(\xi) \phi_k(\xi) \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right]. \end{aligned}$$

Since $Z_n^\xi(t, x, y) = \partial_y u_n^\xi(t, x, y) + s(t, \sigma(\xi))e^x - K$, we derive

$$\begin{aligned} &\mathbb{E} [\phi_k(\xi) \partial_t u_n^\xi] - \mathbb{E} \left[\phi_k(\xi) Z_n^\xi(t, x, y) \left(q_M \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} + q_m \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right) \right] \\ &= \mathbb{E} \left[\phi_k(\xi) (\partial_t u_n^\xi - q_M \partial_y u_n^\xi(t, x, y)) \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right] \\ &\quad + \mathbb{E} \left[\phi_k(\xi) (\partial_t u_n^\xi - q_m \partial_y u_n^\xi(t, x, y)) \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right] \\ &\quad - q_M \mathbb{E} \left[\phi_k(\xi) [s(t, \sigma(\xi))e^x - K] \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right] \\ &\quad - q_m \mathbb{E} \left[\phi_k(\xi) [s(t, \sigma(\xi))e^x - K] \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right] \\ &= \mathbb{E} \left[\phi_k(\xi) D_t^{(-q_M)} u_n^\xi(t, x, y) \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right] + \mathbb{E} \left[\phi_k(\xi) D_t^{(-q_m)} u_n^\xi(t, x, y) \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right] \\ &\quad - q_M W_k^+(t, x, y) - q_m W_k^-(t, x, y) \\ &= \mathbb{E} \left[\phi_k(\xi) D_t^{(-q_m)} u_n^\xi(t, x, y) \right] \\ &\quad + \mathbb{E} \left[\phi_k(\xi) (D_t^{(-q_M)} u_n^\xi(t, x, y) - D_t^{(-q_m)} u_n^\xi(t, x, y)) \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} \right] \\ &\quad - q_M W_k^+(t, x, y) - q_m W_k^-(t, x, y) \\ &= V_k D_t^{(-q_m)} \tilde{u}_k(t, x, y) + \sum_{j=0}^n H_{jk}^+(t, x, y) (D_t^{(-q_M)} \tilde{u}_j(t, x, y) - D_t^{(-q_m)} \tilde{u}_j(t, x, y)) \\ &\quad - q_M W_k^+(t, x, y) - q_m W_k^-(t, x, y). \end{aligned} \tag{2.17}$$

We could also write the equation in terms of $D_t^{(-q_M)} \tilde{u}_k(t, x, y)$. For this, set

$$H_{jk}^-(t, x, y) = \mathbb{E} \left[\phi_j(\xi) \phi_k(\xi) \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right].$$

Then, similar computations give

$$\begin{aligned} \mathbb{E} [\phi_k(\xi) \partial_t u_n^\xi] - \mathbb{E} \left[\phi_k(\xi) Z_n^\xi(t, x, y) \left(q_M \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} + q_m \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right) \right] \\ = V_k D_t^{(-q_M)} \tilde{u}_k(t, x, y) + \sum_{j=0}^n H_{jk}^-(t, x, y) (D_t^{(-q_m)} \tilde{u}_j(t, x, y) - D_t^{(-q_M)} \tilde{u}_j(t, x, y)) \\ - q_M W_k^+(t, x, y) - q_m W_k^-(t, x, y). \end{aligned} \quad (2.18)$$

Interpolating the formulas (2.17) and (2.18) with a parameter $\theta \in \mathbb{R}$ yields

$$\begin{aligned} \mathbb{E} [\phi_k(\xi) \partial_t u_n^\xi] - \mathbb{E} \left[\phi_k(\xi) Z_n^\xi(t, x, y) \left(q_M \mathbf{1}_{Z_n^\xi(t, x, y) \geq 0} + q_m \mathbf{1}_{Z_n^\xi(t, x, y) < 0} \right) \right] \\ = V_k D_t^{(-\theta q_m - (1-\theta)q_M)} \tilde{u}_k(t, x, y) \\ + \sum_{j=0}^n [\theta H_{jk}^+(t, x, y) - (1-\theta) H_{jk}^-(t, x, y)] (D_t^{(-q_M)} \tilde{u}_j(t, x, y) - D_t^{(-q_m)} \tilde{u}_j(t, x, y)) \\ - q_M W_k^+(t, x, y) - q_m W_k^-(t, x, y). \end{aligned} \quad (2.19)$$

▷ **Term with first space derivative.** We simply have

$$\mathbb{E} [\phi_k(\xi) \alpha x \partial_x u_n^\xi(t, x, y)] = \alpha x V_k \partial_x \tilde{u}_k(t, x, y). \quad (2.20)$$

▷ **Term with first second derivative.** The transformation of $\mathbb{E} \left[\phi_k(\xi) \frac{\sigma^2(\xi)}{2} \partial_{xx}^2 u_n^\xi(t, x, y) \right]$ depends on the model of uncertainty on $\xi \mapsto \sigma(\xi)$. In the case **(Q)** where $\sigma^2(\xi) = (\mu + \nu\xi)^2$, in view of the three and four-terms relation (2.11)-(2.12) and using $V_k = \mathbb{E} [\phi_k^2(\xi)] = k!$ (see (2.10)), we have (for $0 \leq k \leq n$)

$$\begin{aligned} \mathbb{E} \left[\phi_k(\xi) \frac{\sigma^2(\xi)}{2} \partial_{xx}^2 u_n^\xi \right] &= \frac{\mu^2}{2} V_k \partial_{xx}^2 \tilde{u}_k + \frac{\nu^2}{2} \sum_{j=0}^n \mathbb{E} [\xi^2 \phi_j(\xi) \phi_k(\xi)] \partial_{xx}^2 \tilde{u}_j \\ &\quad + \mu\nu \sum_{j=0}^n \mathbb{E} [\xi \phi_j(\xi) \phi_k(\xi)] \partial_{xx}^2 \tilde{u}_j, \\ &= \frac{\mu^2}{2} V_k \partial_{xx}^2 \tilde{u}_k + \frac{\nu^2}{2} \left((k+2)(k+1) V_k \partial_{xx}^2 \tilde{u}_{k+2} + (2k+1) V_k \partial_{xx}^2 \tilde{u}_k + V_k \partial_{xx}^2 \tilde{u}_{k-2} \right) \end{aligned}$$

$$\begin{aligned}
& + \mu\nu \left((k+1)V_k \partial_{xx}^2 \tilde{u}_{k+1} + V_k \partial_{xx}^2 \tilde{u}_{k-1} \right) \\
& = V_k \left[\frac{\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k-2} + \mu\nu \partial_{xx}^2 \tilde{u}_{k-1} + \frac{1}{2} (\mu^2 + (2k+1)\nu^2) \partial_{xx}^2 \tilde{u}_k \right. \\
& \quad \left. + \mu\nu(k+1) \partial_{xx}^2 \tilde{u}_{k+1} + \frac{(k+2)(k+1)}{2} \nu^2 \partial_{xx}^2 \tilde{u}_{k+2} \right].
\end{aligned}$$

▷ **Summary.** Now, plug the above into (2.16)-(2.19)-(2.20): after dividing by $V_k = k! > 0$, the PC system in the Hermite basis with the semi-Lagrangian scheme approximated by forward finite differences is ($k = 0, 1, \dots, n$)

$$\begin{aligned}
& - \frac{\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k-2} - \mu\nu \partial_{xx}^2 \tilde{u}_{k-1} - \frac{\mu^2 + (2k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_k \\
& - (k+1)\mu\nu \partial_{xx}^2 \tilde{u}_{k+1} - \frac{(k+2)(k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k+2} + \alpha x \partial_x \tilde{u}_k \\
& + \frac{1}{\Delta t} \tilde{u}_k(t, x, y) - \frac{1}{\Delta t} \tilde{u}_k(t - \Delta t, x, y + [\theta q_m + (1-\theta)q_M]\Delta t) \\
& + \frac{1}{k! \Delta t} \sum_{j=0}^n [\theta H_{jk}^+(t, x, y) - (1-\theta)H_{jk}^-(t, x, y)] \\
& \quad \times (\tilde{u}_j(t - \Delta t, x, y + q_m \Delta t) - \tilde{u}_j(t - \Delta t, x, y + q_M \Delta t)) \\
& - \frac{q_m}{k!} W_k^-(t, x, y) - \frac{q_M}{k!} W_k^+(t, x, y) = 0.
\end{aligned}$$

3 Implementation

Experimentally, it turns out that the interpolation parameter θ in (2.19) does not noticeably affect the results. Thus, we fix $\theta = 1$ henceforth. The PC system reads:

$$\begin{aligned}
& - \frac{\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k-2} - \mu\nu \partial_{xx}^2 \tilde{u}_{k-1} - \frac{\mu^2 + (2k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_k \\
& - (k+1)\mu\nu \partial_{xx}^2 \tilde{u}_{k+1} - \frac{(k+2)(k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k+2} + \alpha x \partial_x \tilde{u}_k \\
& + \frac{1}{\Delta t} \tilde{u}_k(t, x, y) - \frac{1}{\Delta t} \tilde{u}_k(t - \Delta t, x, y + q_m \Delta t) \\
& + \frac{1}{k! \Delta t} \sum_{j=0}^n H_{jk}^+(t, x, y) (\tilde{u}_j(t - \Delta t, x, y + q_m \Delta t) - \tilde{u}_j(t - \Delta t, x, y + q_M \Delta t)) \\
& - \frac{q_m}{k!} W_k^-(t, x, y) - \frac{q_M}{k!} W_k^+(t, x, y) = B_k(t, x, y), \tag{3.1}
\end{aligned}$$

where for convenience we have introduced $B_k(t, x, y) = 0$ (it will be nonzero with the benchmark problem to be introduced in Section 4). The above system poses the challenge that the nonlinear terms $\{H_{ij}^+, W_k^\pm\}$ couple all of the components of the solution, $\tilde{u}_0, \dots, \tilde{u}_n$. This means that the discretized version of (3.1) with finite differences (FD) will in principle have full bandwidth. Indeed, a well known drawback of PC is that the method may give rise to taxing algebraic equations [Xiu09], which is only compounded by the nonlinear nature of (3.1).

In order to (partially) circumvent this issue, we put forward an FD semi-implicit time-stepping scheme, whereby the linear terms are evaluated implicitly, and the nonlinear ones, explicitly (i.e. based on the solution on the previous timestep). Formally:

$$\begin{aligned}
& -\frac{\nu^2}{2} \left[\partial_{xx}^2 \tilde{u}_{k-2} \right]_{FD}^{m+1} - \mu\nu \left[\partial_{xx}^2 \tilde{u}_{k-1} \right]_{FD}^{m+1} - \frac{\mu^2 + (2k+1)\nu^2}{2} \left[\partial_{xx}^2 \tilde{u}_k \right]_{FD}^{m+1} \\
& - (k+1)\mu\nu \left[\partial_{xx}^2 \tilde{u}_{k+1} \right]_{FD}^{m+1} - \frac{(k+2)(k+1)\nu^2}{2} \left[\partial_{xx}^2 \tilde{u}_{k+2} \right]_{FD}^{m+1} + \left[\alpha x \partial_x \tilde{u}_k \right]_{FD}^{m+1} \\
& + \frac{1}{\Delta t} \left[\tilde{u}_k(t, x, y) \right]_{FD}^{m+1} - \frac{1}{\Delta t} \left[\tilde{u}_k(t - \Delta t, x, y + q_m \Delta t) \right]_{FD}^{m+1} + \frac{1}{k! \Delta t} \sum_{j=0}^n \left[H_{jk}^+(t, x, y) \right]_{FD}^m \\
& \quad \times \left(\left[\tilde{u}_j(t - \Delta t, x, y + q_m \Delta t) \right]_{FD}^m - \left[\tilde{u}_j(t - \Delta t, x, y + q_M \Delta t) \right]_{FD}^m \right) \\
& - \frac{q_m}{k!} \left[W_k^-(t, x, y) \right]_{FD}^m - \frac{q_M}{k!} \left[W_k^+(t, x, y) \right]_{FD}^m = B_k^{m+1}.
\end{aligned} \tag{3.2}$$

In (3.2), $[\mathcal{L}v]_{FD}^{m+1}$ and $[\mathcal{L}v]_{FD}^m$ are proper FD discretizations of some operator \mathcal{L} evaluated at current time $t^{m+1} = (m+1)\Delta t$ and at previous time $t^m = m\Delta t$, respectively. In this way, the discretized system (3.2) has a pentadiagonal, rather than full, coupling. Let the spatial domain $[x_m = -x_L, x_M = x_L] \times [0, y_M]$ be discretized into a grid of $(1+n_x)(1+n_y)$ nodes $(x_i, y_j) = (-x_L + i\Delta x, j\Delta y)$, where $\Delta x = 2x_L/n_x$ and $\Delta y = y_M/n_y$; and let $[U_k]_{ij}^{m+1}$ be the FD approximation to $\tilde{u}_k(t^{m+1}, x_i, y_j)$. Furthermore, the random dimension is approximated as $\Omega \approx [-L_w, L_w]$ and discretized into $n_w + 1$ equispaced nodes. We have employed the following FD approximations:

- central second-order—i.e. $\mathcal{O}(\Delta^2 x)$ —FDs for the implicit second derivative,

$$\left[\partial_{xx}^2 \tilde{u}_k(t^{m+1}, x_i, y_j) \right]_{FD}^{m+1} = \frac{[U_k]_{i+1,j}^{m+1} + [U_k]_{i-1,j}^{m+1} - 2[U_k]_{i,j}^{m+1}}{(\Delta x)^2};$$

- second-order upwind FDs for the implicit first derivative along x ,

$$\left[c \partial_x \tilde{u}_k(t^{m+1}, x_i, y_j) \right]_{FD}^{m+1} = \max(c, 0) \frac{3[U_k]_{ij}^{m+1} - 4[U_k]_{i-1,j}^{m+1} + [U_k]_{i-2,j}^{m+1}}{2\Delta x}$$

$$+ \min(c, 0) \frac{-[U_k]_{i+2,j}^{m+1} + 4[U_k]_{i+1,j}^{m+1} - 3[U_k]_{ij}^{m+1}}{2\Delta x};$$

- second-order central FDs for the explicit first derivative along y (notice that there is no need for an upwind scheme since it is computed explicitly),

$$\left[Z(t^m, x_i, y_j, \xi) \right]_{FD}^m = \frac{\sum_{k=0}^n ([U_k]_{i,j+1}^m - [U_k]_{i,j-1}^m) \phi_k(\xi)}{2\Delta y} + s(t^m, \sigma(\xi)) e^{x_i} - K.$$

The nonlinear terms are approximated by numerical integration. Since they involve the indicator function, straightforward trapezoidal quadrature is preferable to higher-order methods, which are more prone to oscillations when dealing with discontinuous integrands (due to the indicator). For example,

$$\begin{aligned} \left[H_{sk}^+((t^{m+1}, x_i, y_j)) \right]_{FD}^m &= \frac{L_w}{n_w} \left(\phi_s(-w_L) \phi_k(-w_L) p(-w_L) \mathbf{1}_{[Z(t^m, x_i, y_j, -w_L)]^m} \right. \\ &\quad \left. + \phi_s(w_L) \phi_k(w_L) p(w_L) \mathbf{1}_{[Z(t^m, x_i, y_j, w_L)]^m} + 2 \sum_{r=1}^{n_w} \phi_s(\xi_r) \phi_k(\xi_r) p(\xi_r) \mathbf{1}_{[Z(t^m, x_i, y_j, \xi_r)]^m} \right). \end{aligned}$$

When choosing the timestep, we mind $\Delta t < q_M \Delta y$ —i.e. prevent the interpolation of semilagrangian derivatives from “jumping” over cells. (A CFL upper bound $\Delta t < 2\Delta x/|\alpha x_L|$ is not necessary since the derivatives in x are evaluated implicitly.)

The interpolation scheme is a critical aspect of semilagrangian differentiation—specially so if the former is to be monotone, see [War16]. In particular, the condition

$$\frac{1}{\Delta t} \left[\tilde{u}_k(t - \Delta t, x, y + q_m \Delta t) \right]_{FD}^{m+1} \rightarrow 0 \quad \text{as} \quad \Delta t \rightarrow 0$$

must hold for consistency. (In our Matlab implementation, this was achieved with cubic-spline interpolation—but not necessarily so with the other available methods.)

The resulting linear system to be solved at each time step has a block pentadiagonal structure (there are linear algebra routines tailored to such systems, see [BF07]).

Hamilton-Jacobi PDEs such as (2.5) are well known to often lack a solution in the classical sense [BS91]. For that reason, numerical schemes for solving scalar equations such as (2.5) must be designed to pick up the unique, sought-for viscosity solution. With FDs in particular, this is enforced with monotone discrete operators [BS91, Tou13]. In contrast to the scalar case—and to the best of our knowledge—there is not a systematic theory of viscosity solutions for systems of equations such as (3.1) [LYL03]. Therefore, we will always assume that either a classical solution does exist, or that our numerical method is able to pick up the viscosity solution, otherwise. With this proviso, we have

dropped the requirement of monotonicity of the finite difference operators. For instance, (3.3) and cubic-spline interpolation are not monotone, but they are more accurate than monotone versions, which typically are just $\mathcal{O}(\Delta x)$.

The code has been written in vectorized style in Matlab R2017 and run on a laptop.

4 Numerical validation

In order to assess empirically the accuracy and stability of the numerical scheme, we consider the PDE in (2.5) with an additional source term B :

$$\partial_t u - \frac{\sigma^2(\xi)}{2} \partial_{xx}^2 u + \alpha x \partial_x u - Z(q_M \mathbf{1}_{Z>0} + q_m \mathbf{1}_{Z<0}) = B(t, x, y, \xi). \quad (4.1)$$

Analogously to Section 2.2.3, we derive the PC equations for (4.1) up to truncation level n . For convenience, let us rewrite the resulting PC system ($k = 0, \dots, n$) without the semilagrangian approximation as

$$\begin{aligned} & \partial_t \tilde{u}_k - \frac{\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k-2} - \mu \nu \partial_{xx}^2 \tilde{u}_{k-1} - \frac{\mu^2 + (2k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_k - (k+1)\mu \nu \partial_{xx}^2 \tilde{u}_{k+1} \\ & - \frac{(k+2)(k+1)\nu^2}{2} \partial_{xx}^2 \tilde{u}_{k+2} + \alpha x \partial_x \tilde{u}_k - \frac{1}{k!} \mathbb{E} \left[\phi_k(\xi) Z(t, x, y, u_n(\xi)) (q_M \mathbf{1}_{Z>0} + q_m \mathbf{1}_{Z<0}) \right] \\ & = \frac{1}{k!} \mathbb{E}[\phi_k(\xi) B(t, x, y, \xi)]. \end{aligned} \quad (4.2)$$

The difference between (4.2) and (2.16) is the presence of the vector of source terms $\mathbb{E}[\phi_k(\xi) B(t, x, y, \xi)]/k!$ ($k = 0, \dots, n$).

Let us choose as exact solution of the benchmark problem (or simply, the benchmark)

$$u_{ex}(t, x, y, \xi) = \sum_{k=0}^{n_{ex}} \tilde{v}_k(t, x, y) \phi_k(\xi), \quad (4.3)$$

with

$$\tilde{v}_k(t, x, y) = \frac{1+t}{(k+1)\sqrt{k!}} \sin\left(\frac{2\pi x}{3}\right) \sin\left(\frac{\pi y}{2}\right),$$

which complies with (2.14). Moreover, we let $u_{ex}(t, x, y, \xi)$ serve as initial condition and Dirichlet BCs for the benchmark problem (4.1). (Note that it can no longer be associated with the fair price of a gas swing option).

We highlight from the onset that results turn out to be remarkably insensitive to discretization in the random dimension (provided it is adequate). Therefore, we set

$L_w = 13$, $n_w = 200$ in all the ensuing experiments. Regarding the spatial domain, we set $x_m = -3$, $x_M = 3$, $y_M = 4$, and take Dirichlet BCs along $\{|x| = 3, 0 \leq y \leq y_M\}$ and $\{x_m \leq x \leq x_M, y = y_M\}$. The spatial discretization is discussed in terms of either n_x and n_y , or of $\Delta x = 6/n_x$ and $\Delta y = 4/n_y$, equally. The remaining parameters are: $F(0, t) = .15(1 + (\sin 2\pi t)/3)$, $K = .1$, $q_m = .1$, $q_M = .7$, $\mu = 1$, $\alpha = .3$ and (unless stated otherwise) $\nu = .3$. The variables n and T remain to be set, as well as the value n_{ex} for the u_{ex} in (4.3). Given $N_{ev} = 500$ fixed, uniformly scattered points across the domain, the error of the PC expansion is characterized by

$$RMS(t) = \sqrt{\frac{1}{N_{ev}} \sum_{j=1}^{N_{ev}} \sum_{k=0}^n (\tilde{v}_k(t, x_j, y_j) - \tilde{u}_k(t, x_j, y_j))^2}.$$




Figure 1: Convergence of $RMS(t = 1)$ w.r.t. Δt , Δx and Δy (separately), with $n = n_{ex}$.

In a first set of experiments, we enforce the truncation error of the PC expansion to be zero by always setting $n_{ex} = n$, and focus on the discretization error with respect to Δt , Δx and Δy . This is shown both with $n = n_{ex} = 0$ (left side of Figure 1), and with $n = n_{ex} = 16$ (right side). A few remarks are in place:

- Using cubic-spline interpolation, consistency with respect to Δt is clear (also for smaller values than shown on Figure 1). Furthermore, stability is apparent.
- While the accuracy noticeably improves with decreasing Δt and Δx , it is barely sensitive to Δy .
- The pentadiagonal coupling notwithstanding, accuracy is scarcely affected by the number $n + 1$ of PC equations in the system. However, when n is very large, the evaluation of high-order Hermite polynomials may overflow (not shown). This is a well known issue which can be ameliorated by ad-hoc procedures, for instance by resorting to the Cauchy-Schwartz bound

$$|H_{jk}^{\pm}(t, x, y)| \leq \mathbb{E}[|\phi_j(\xi)| |\phi_k(\xi)|] \leq \sqrt{j!} \sqrt{k!}.$$

Then, if the numerical evaluation of $|H_{jk}^{\pm}(t, x, y)|$ overshoots the above bound, $H_{jk}^{\pm}(t, x, y)$ must be replaced by $\text{sign}(H_{jk}^{\pm}(t, x, y)) \sqrt{j!} \sqrt{k!}$.

- With $n = 0$, the PC system is effectively an HJB equation. While the FD operators being used are not monotone, there is no need for them in the first place, since the benchmark solution is guaranteed to be smooth.

In Figure 1, the RMS error decreases roughly linearly w.r.t. Δt as $n_x \rightarrow \infty$ and $n_y \rightarrow \infty$. In order to better assess the convergence rate of the numerical scheme, we produce an ensemble of simulations with discretization values randomly picked from uniform distributions in the ranges $0 \leq n = n_{ex} \leq 20$, $50 \leq \Delta x \leq 200$, $50 \leq \Delta y \leq 200$, and $\log(10^{-3}) \leq \log(\Delta t) \leq \log(.025)$. Then, we fit the resulting $RMS(t = 1)$ errors to the nonlinear model

$$RMS(t = 1) = (1 + c_4 n) (c_1 \Delta t + c_3 (\Delta x)^{c_2}),$$

yielding as best-fit constants $c_1 \approx .26$, $c_2 \approx 2.26$, $c_3 \approx .55$ and $c_4 \approx .0025$, roughly in agreement with the expected second order convergence rate along x . This fit is plotted in Figure 2, and it turns out to be distinctly the best (in a least-squares sense) over several plausible nonlinear convergence models, including floating powers of Δt and Δy (not shown). (The outliers in Figure 2 are simulations with small Δt and large Δx .)

In the second part of this numerical study we focus on the effect of the truncation level n . Henceforth, we set $n_{ex} = 16$ and allow n to differ from it.

In Figure 3, the longer-term behaviour of the RMS error is plotted. Observe that:

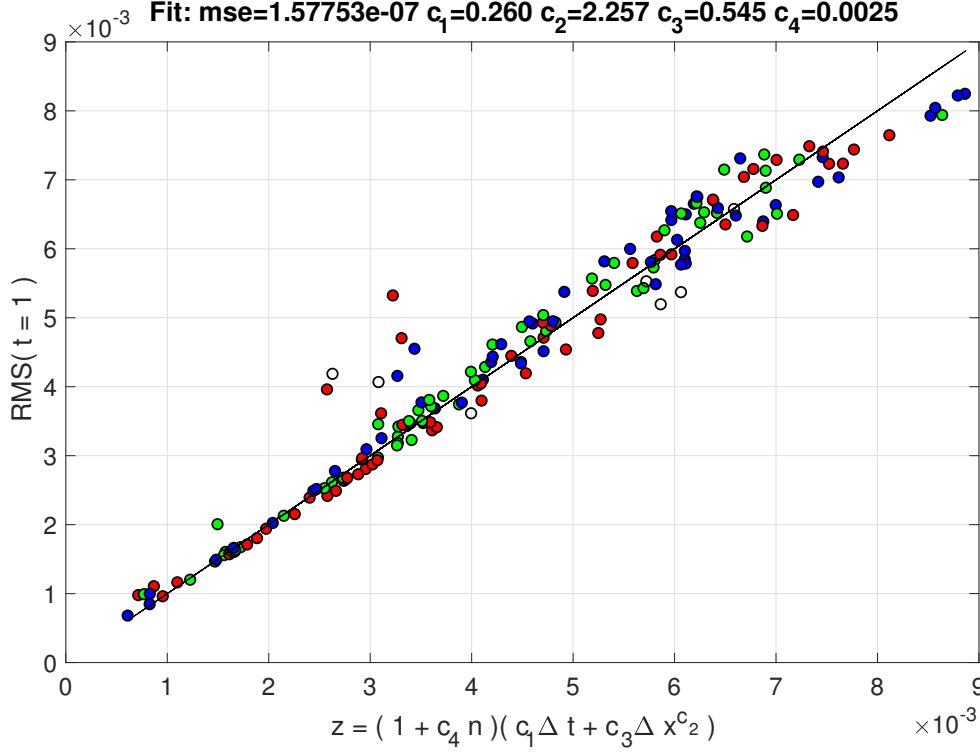


Figure 2: Best fit of $RMS(t=1)$ error to a nonlinear convergence rate model. Colours: $n=0$ (white), $1 \leq n \leq 6$ (green), $7 \leq n \leq 13$ (blue), $14 \leq n \leq 20$ (red); and $n_{ex} = n$.

- Asymptotically, all the instances of the RMS error tend to grow at a constant, sublinear rate (check the dashed lines of slope one on Figure 3). This means that relative errors are actually decreasing with increasing t (recall that the benchmark solution (4.3) grows linearly with t). Moreover, it strongly suggests stability of the numerical PC code for longer terminal times.
- When the PC solution lives in a smaller random space than the benchmark (i.e. when $n < n_{ex} = 16$), errors are almost always larger when ν is larger. On the other hand, when $n = n_{ex} = 16$, the RMS error is virtually insensitive to ν . Moreover, truncating beyond n_{ex} does not improve accuracy (as expected).

A well known issue of PC is that the moments of the PC solution (specially from the variance onwards) tend to deteriorate rather rapidly as $t \rightarrow \infty$, unless the truncation level is also constantly raised or the random basis adaptively changed [GvVK10]. Figure 4 shows the propagation of the error of the mean and that of the variance (with respect

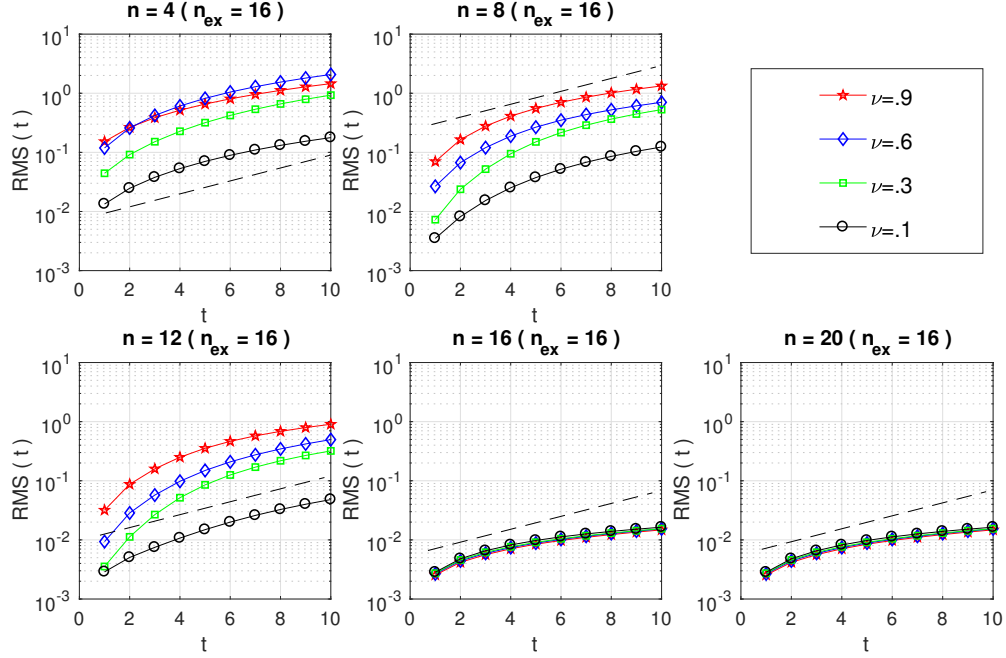


Figure 3: Propagation of the RMS error. Auxiliary dashed lines have slope 1.

to the mean and variance of the benchmark solution) at $(x_0, y_0) = (\sqrt{2}, \sqrt{3})$. The relative error of the variance grows faster than that of the expectation and—unless n is very close to n_{ex} —becomes $\mathcal{O}(1)$ by time $t \gtrsim 2$.

Finally, in Figure 5 we study the dependence of the PC numerical solution on the random variable $\xi \sim \mathcal{N}(0, 1)$. We focus on the interval $-3 \leq \xi \leq 3$ which has an aggregate probability of over 99%. When $n \lesssim n_{ex} = 16$, the error (with respect to the benchmark) is small and increases with T and with $(\Delta t, \Delta x, \Delta y)$. On the other hand, when n is too low, the truncation error overrides that of discretization and can be unacceptably large ($\mathcal{O}(1)$ and bigger).

Summing up, the numerical scheme (3.2) has been validated with a smooth solution.

5 Uncertainty quantification on gas swing option

We consider the gas swing option priced under one-factor dynamics in [BBP09]. The parameters are: $T = 1$ (in years), $\mu = .7$, $\alpha = 4$, $F_{0t} = 20$ (a constant), $q_m = 0$, $q_M = 2190$, $Q_{max} = 1900$, $Q_{min} = 1300$ and $K = \{5, 10, 15, 20\}$. In a slight departure from our own framework, the contract described there does not allow to adjust the gas




Figure 4: Error of the expectation (left) and variance (right) of the PC expansion for various truncation levels ($\Delta t = .005$, $n_x = 200$, $n_y = 100$). Dashed lines have slope 1.

drawing rate continuously, but only once a day over one year (i.e. it depends on 365 discrete variables rather than on a continuous function $y(t)$). There are two cases:

- Unconstrained contract: in this scenario, there are no upper or lower limits to the total amount of gas that can be withdrawn over the entire contract period (the swing option contract is in fact equivalent to a string of 365 call options). Therefore, $P(x, y) = 0$ in (2.5) (so that the initial condition is smooth).
- Constrained contract: this is the most usual contractual scenario, when there are strict limits on the total amount of gas that can be withdrawn, i.e.

$$Q_{min} \leq Y_T \leq Q_{max}. \quad (5.1)$$

Strictly speaking, (5.1) amounts to a penalty function which is zero for $Q_{min} \leq y \leq Q_{max}$ and infinity outside that interval. For numerical purposes, it is approximated by the double ramp function

$$P(x, y) = A s(0, \sigma) e^x (\min(0, y - Q_{min}) - \max(0, y - Q_{max})),$$




Figure 5: Relative error in the random dimension with respect to the benchmark ($n_{ex} = 16$) at $(x_0, y_0) = (\sqrt{2}, \sqrt{3})$ and: $t = 1$ (left), $t = 5$ (middle), and $t = 10$ (right). Dashed curves represent a rougher discretization ($\Delta t = .02, n_x = 90, n_y = 60$). Solid curves are a finer discretization ($\Delta t = .005, n_x = 150, n_y = 100$). For $n = 4$, both sets are nearly indistinguishable to the naked eye.

where A is a user-defined positive constant. We take $A = 10000$ as recommended by [BBP09]. Note that $A = 0$ recovers the unconstrained case.

The solution is to be evaluated at $y_T = 0$ (no initial gas) and

$$x_T = \frac{\mu^2}{4\alpha}(1 - e^{-2\alpha T}) + \log \frac{F_{0t}(T)}{F_{0t}(0)} \approx .03061.$$

The initial conditions for the PC equations now read:

$$\tilde{u}_k(0, x, y) = \frac{Ae^x}{k!} (\min(0, y - Q_{min}) - \max(0, y - Q_{max})) \mathbb{E} \left[s(\sigma(\xi), 0) \phi_k(\xi) \right], \quad (5.2)$$

which is again approximated by trapezoidal quadrature for simplicity. Numerical experiments show that taking $L_w = 10$ and $n_w = 100$ is sufficient in all the cases.

Due to the nonsmooth initial condition, in this problem linear interpolation performs better than cubic splines, so that we use the former instead.




Figure 6: Illustration of the “boundary pollution” effect with $K = 5$ and $n = 8$: $x_M \gtrsim 15\mu/\sqrt{2\alpha}$ for the solution to be unaffected by the finite computational domain.

Computational domain and BCs. As discussed in Section 2.1, the computational domain $D = [x_m, x_M] \times [0, y_M > q_M T]$ must be large enough that the numerical solution at point (x_T, y_T) at time T is not affected by the BCs enforced up to time T . We set $y_M = 1.3q_M T = 2847$, and choose as BCs:

$$u_n^{\mu+\nu\xi}(t, x, y) = \begin{cases} P(x, y), & \text{if } (x, y) \in \partial D \text{ and } t = 0, \\ 0, & \text{if } (x, y) \in \partial D \text{ and } 0 < t \leq T. \end{cases}$$

In this way, only the effect of the initial condition along ∂D must be reckoned with. Owing to the e^x factor in (5.2), $u_n^{\mu+\nu\xi}(T, x_T, y_T)$ is rather insensitive to $x_m < 0$ (we set $x_m = -1.24$), but $x_M > 0$ must be handled with more care. The reason is that the influence of the initial condition on the boundary decays as $\mathcal{O}(e^{-x_M^2/(2\mu^2)})$, but the BC itself grows as $\mathcal{O}(e^{x_M})$. The value of x_M is estimated as follows: we wish to keep the “boundary pollution” (let us call it b) smaller than 0.1% of the solution at (T, x_T, y_T) . According to Table 1 (more on it later), $b = 30 > .001 \times \mathbb{E}[u_n^{\mu+\nu\xi}(T, x_T, y_T)] \geq .001 \times |\tilde{u}_k(T, x_T, y_T)|$ ($0 \leq k \leq n$) is a conservative bound. Combining the propagation estimate (2.6) and the BC (at $t = 0$) given by (5.2) yields $x_M \approx 15\mu/\sqrt{2\alpha} \approx 3.72$ (i.e. 15 standard deviations of the OU process (2.1)). Figure 6 shows the effect of the “boundary pollution” for several x_M .

	unconstrained ($A = 0$)		constrained ($A = 10000$)	
K	$u_0^\mu(T, x_T, y_T)$	u_{EX}	$u_0^\mu(T, x_T, y_T)$	u_{RQ}
5	33274	32760	29361	29342
10	22298	21844	19987	19866
15	11760	11381	10867	10698
20	4181	3966	2580	2680

Table 1: Using the PC algorithm (3.2) with $n = 0$ and $\nu = 0$ as an HJB solver for pricing gas swing options without uncertainty. ($\Delta t = 1/365, n_x = 320, n_y = 200, x_M = 3.72$). References u_{EX} and u_{RQ} are the solution of the slightly different problem in [BBP09].

Using the PC code in “scalar mode”. The PC algorithm can, in principle, also be used for scalar HJB equations such as (2.5)—and hence for pricing swing options without uncertainty, too. This is possible by letting $n = 0$, $\nu = 0$, and $\mu = \sigma$ in (3.1). In the sequel, $u_0^\sigma(\cdot)$ is to be understood as the numerical solution of an HJB equation with volatility σ obtained by so using the PC code (in “scalar mode”).

Since some of the FD operators in Section 3 are not monotone, $u_0^\sigma(\cdot)$ is not guaranteed to pick up the viscosity solution. (This is particularly relevant in the constrained case, where the initial conditions are non-differentiable.) On the other hand, comparison with the results of a separate viscosity-capturing HJB code (based on Howard’s algorithm [Tou13]) shows no difference with $u_0^\sigma(\cdot)$ —while the latter enjoys a higher convergence rate w.r.t. Δx . Therefore, we stick to the PC code henceforth.

In order to get a feeling about the several discretization parameters involved, we calculate the swing option prices for $K = \{5, 10, 15, 20\}$ with no uncertainty (i.e. $u_0^\sigma(T, x_T, y_T)$), and compare it with the values given in [BBP09]. In Table 1, u_{EX} has been calculated with the Black-Scholes formula, and u_{RQ} with a quantization method (see [BBP09] for details). In the remaining of the paper, we will use either the

- “coarser discretization”: $\Delta t = 1/365, n_x = 160$ (and $n_y = 200$), or the
- “finer discretization”: $\Delta t = 1/(2 \times 365), n_x = 320$ (and $n_y = 200$ as well).

Reference solution. In order to assess the quality of the PC solution, we will fix ν first, and then compare $u_n^{\mu+\nu\xi}(T, x_T, y_T)$ inside a given interval of ξ with the HJB solution $u_0^\sigma(T, x_T, y_T)$, where $\sigma = \mu + \nu\xi$. Specifically, for each K we precompute a set of $\{u_0^{\sigma_i}(T, x_T, y_T)\}_{i=1}^{90}$ HJB solutions with the “finer” discretization at 90 equispaced volatilities $\sigma_1 = 0.1, \dots, \sigma_{90} \approx 1.75$. For a given ν , the comparable interval of ξ is thus $(\sigma_1 - \mu)/\nu \leq \xi \leq (\sigma_{90} - \mu)/\nu$.




Figure 7: Comparison between the PC expansion with increasing n , and the HJB reference curves, at (T, x_T, y_T) . With $n > 4$, the improvement becomes harder to tell with the naked eye; use the zoom and/or check Table 2. ($\nu = .5$; “finer” discretization)

We start the numerical study by investigating the interplay between the discretization and the truncation level, n . In order to better highlight the effect of n , we first pick the relatively large value $\nu = .5$. The dashed curves in Figure 7 are based on a subset of 70 of those HJB simulations. They represent $u_0^\sigma(T, x_T, y_T)$ in the interval $0.1 \leq \sigma \leq 1.3$ (i.e. centered around μ). It can be seen that the delta of the swing option is nearly constant.

The red curves in Figure 7 are the PC approximation $u_n^{\mu+\nu\xi}(T, x_T, y_T)$ for $n = \{2, 4, 8\}$. The fit to the reference HJB curves improves as n grows. In order to systematically check this, we track the RMS error between each pair of curves, defined as (we

K	$n = 2$		$n = 4$		$n = 8$		$n = 16$	
5	3303.3	3289.6	1563.0	1621.8	1083.7	1151.5	—	—
10	1834.9	1791.3	686.9	630.0	536.9	432.8	565.3	479.6
15	1200.3	1051.7	854.9	490.6	837.0	443.0	806.7	414.1
20	2647.7	2261.0	2255.3	1782.3	1872.7	1323.4	1660.3	1077.2
disc.	coarser	finer	coarser	finer	coarser	finer	coarser	finer

Table 2: Improvement of the PC solution as n grows and the discretization is refined, as measured by R^K (5.4). For $K = 5$ and $n = 16$, both discretizations are inadequate.

include a superindex K to specify the pair of curves):

$$R^K = \sqrt{\frac{1}{70} \sum_{i=1}^{70} \left(u_0^{K, \sigma_i}(T, x_T, y_T) - u_n^{K, \sigma_i}(T, x_T, y_T) \right)^2}. \quad (5.4)$$

The values of R^K for the “finer” and “coarser” discretization are listed on Table 2. As expected, R^K usually improves (i.e. diminishes) as the discretization is refined and n grows. The exception are the entries with $K = 5$, for which neither discretization is fine enough at $\nu = .50$. (Since Hermite polynomials oscillate with ever larger amplitude as $|\xi|$ increases, numerical errors accumulate at the tails of the higher-order components of $u_n^{\mu+\nu\xi}(\cdot)$ —check also Figure 5.)

Figure 8 zooms in on the fit between $u_n^{\mu+\nu\xi}(T, x_T, y_T)$ and the HJB reference curve with $K = 15$, for $n = \{2, 4, 8, 16\}$ and the “finer” discretization. It is apparent that as n grows, the errors accumulate on the higher end of ξ . Therefore, the truncation level can be increased only so much for a given discretization. For this reason, when computing the cdf of the swing prices under uncertainty, the integration w.r.t. ξ must be cut short above. Furthermore, since $\sigma = \mu + \nu\xi$ must be positive, swing prices for $\xi < -\mu/\nu$ are meaningless. In sum, we introduce the following truncated cdf:

$$\text{CDFT}[u_n^{\mu+\nu\xi}, z] := \frac{1}{\sqrt{2\pi}} \int_{\max(\xi_m, -\mu/\nu)}^{\xi_M} e^{-\xi^2/2} \mathbf{1}_{u_n^{\mu+\nu\xi}(T, x_T, y_T) \leq z} d\xi,$$

where ξ_m and ξ_M are respectively negative and positive enough that the distribution is well represented. (Note that this may not be possible if the cut-off $|\mu/\nu|$ is too small).

By differentiating $\text{CDFT}[u_n^{\mu+\nu\xi}, z]$ with respect to z , one obtains the approximate pdf of the Swing option price under uncertainty. The error tends to accumulate at the tails, thus having a larger adverse effect on the higher moments of the distribution. This is illustrated in Figure 9, where the more realistic uncertainty $\nu = .1$ has been




Figure 8: Comparison of the PC expansions (“finer” discretization) with increasing n to the reference HJB curve for $K = 15$ and $\nu = .5$. For the sake of clarity, only a few individual values of the PC curves are displayed instead of the continuous curves.

used. In that Figure, the PC curves with different n (but the same discretization) are nearly indistinguishable. The black curve, based on $CDFT[u_0^\sigma, z]$, is equivalent to a Monte Carlo reference pdf, but without the statistical error.

In Figure 9, and in all the remaining simulations in the paper, the following values have been used: $K = 15$, $\xi_m = -3.5$, $\xi_M = 3.5$. (We point out that no smoothing has been applied either to the truncated cdfs or to the pdfs.)

6 Uncertainty Value Adjustment

Consider now the following practical situation. The swing option is sold at the price $v^\sigma(0)$ corresponding to a medium estimation of the volatility: this may be an equilibrium price in the market. However, the seller of the contract wants to assess the




Figure 9: Pdf of the swing option price (with $K = 15$) under 10% uncertainty on the volatility. The error of discretization overrides that of truncation. The black curve is the reference “Monte Carlo” pdf, based on stochastic collocation of the HJB equation.

risk for volatility misspecification: he could consider a worst-case analysis by computing $\sup_{\sigma} v^{\sigma}(\sigma)$ but this may be too conservative. More reasonably, he incorporates the distribution on σ (coming from his calibration/estimation procedure) and aims at computing $\mathbb{E}[(v^{\sigma}(\sigma) - z)_{+}]$ with $z \geq v^{\sigma}(0)$. The above quantity can be seen as an Uncertainty Value Adjustment (UVA), whose purpose is to give a financial evaluation of the possibility of misspecification of σ in the swing problem, beyond a risk threshold z . In finance (like for CVaR), z is given by a quantile at 95% or 99%.

We proceed to explore this idea using the gas swing option in the previous Section. Let us consider three quantiles $q = \{99\%, 95\%, 90\%\}$. For each of them, and given a PC expansion $u_n^{\mu+\nu\xi}(T, x_T, y_T)$, the risk threshold z_q is defined by

$$\text{CDFT}[u_n^{\mu+\nu\xi}, z_q] = q/100.$$

After solving for z_q with Newton’s method, the UVA is calculated as

$$\text{UVA}[u_n^{\mu+\nu\xi}, q] = \frac{1}{\sqrt{2\pi}} \int_{\max(\xi_m, -\mu/\nu)}^{\xi_M} \min(0, u_n^{\mu+\nu\xi} - z_q) e^{-\xi^2/2} d\xi.$$

Table 3 compiles UVA approximations for $K = 15$ and $\nu = 10\%$ using several PC expansions. The entries are the percentage of UVA relative to the mean price of the

$\nu = .10$ and $K = 15$

q	MC	$n = 2$		$n = 4$		$n = 8$		$n = 16$	
99%	.0067	.0072	.0070	.0071	.0069	.0071	.0069	.0071	.0069
95%	.0453	.0476	.0463	.0470	.0457	.0470	.0457	.0470	.0457
90%	.1031	.1075	.1045	.1065	.1036	.1066	.1035	.1066	.1036
disc.		coarser	finer	coarser	finer	coarser	finer	coarser	finer

Table 3: Relative (percentual) UVA: $(100 \times \text{UVA}[u_n^{\mu+\nu\xi}, q] / \mathbb{E}[u_n^{\mu+\nu\xi}(T, x_T, y_T)])\%$. For the Monte Carlo (MC) references, the entries are $(100 \times \text{UVA}[u_0^\sigma, q] / \mathbb{E}[u_0^\sigma(T, x_T, y_T)])\%$.

swing option. They are compared with the “Monte Carlo” reference based on stochastic collocation of the HJB equation. A relatively coarse discretization and low n are enough to get adequate accuracy.

With $\mu = .7$ and $\nu = .25$ the cut-off value $|\mu/\nu| = 2.8$ is already small (less than 3 standard deviations of the distribution of ξ). In order to investigate UVA at larger uncertainties than 10%, we set $\mu = 1.3$ (and keep $K = 15$, $\xi_m = -3.5$, $\xi_M = 3.5$, and the other parameters as before). See Table 4 for the UVA percentages in that case.

ν	$q = 99\%$		$q = 95\%$		$q = 90\%$	
10%	.0062	.0062	.0424	.0424	.0973	.0974
20%	.0126	.0128	.0857	.0858	.1964	.1960
30%	.0197	.0204	.1322	.1323	.3014	.2990
40%	.0267	.0276	.1832	.1797	.4156	.4042
n	2	4	2	4	2	4

Table 4: UVA (in terms of a percentage over the expected price of swing option) for $\mu = 1.3$, $K = 15$, and the other parameters as in Section 5. (The “coarse” discretization was employed.)

We conclude that settling for $n = 2$ and the “coarser” discretization is adequate for an accurate estimation of UVA. On a laptop and in Matlab, it takes about 3–4 minutes to calculate $u_2^{\mu+\nu\xi}(\cdot)$ and the UVA. While the risk premium against misspecification of the volatility is quite small when the calibration is precise, it grows roughly linearly as the calibration uncertainty increases.

References

- [Bas15] Basel Committee on Banking Supervision. Review of the credit valuation adjustment risk framework. *Bank for International Settlements*, July 2015.
- [BBP09] O. Bardou, S. Bouthemy, and G. Pages. Optimal quantization for the pricing of swing options. *Applied Mathematical Finance*, 16(2):183–217, 2009.
- [BEBD⁺06] 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. *Methodology and Computing in Applied Probability*, 8(4):517–540, 2006.
- [BF07] K. Benkert and R. Fischer. An efficient implementation of the Thomas algorithm for block pentadiagonal systems on vector computers. In: *Shi, Y., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) Proceedings of the 7th International Conference on Computer Science, ICCS*, pages 144–151, 2007.
- [BL14] P. Briand and C. Labart. Simulation of BSDEs by Wiener Chaos Expansion. *Annals of Applied Probability*, 24(3):1129–1171, 2014.
- [Bre91] M. J. Brennan. The price of convenience and the valuation of commodity contingent claims. *Stochastic Models and Option Values*, 200(22–71), 1991.
- [BS91] G. Barles and P. E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptot. Anal.*, 4:271–283, 1991.
- [CSK01] L. Clewlow, C. Strickland, and V. Kaminski. Valuation of swing contracts in trees. *Energy and Power Risk Management*, 6(4):33–34, 2001.
- [Gob02] E. Gobet. LAN property for ergodic diffusion with discrete observations. *Ann. Inst. H. Poincaré Probab. Statist.*, 38(5):711–737, 2002.
- [GvVK10] M. Gerristma, J. B. van der Steen, P. Vos, and G. E. Karniadakis. Time-dependent generalized polynomial chaos. *J. Comput. Phys.*, 229(22):8333–8363, 2010.
- [Har89] A.C. Harvey. *Forecasting, Structural Time Series Analysis, and the Kalman Filter*. Cambridge, 1989.

- [HS14] T. Huschto and S. Sager. Solving stochastic optimal control problems by a Wiener chaos approach. *Vietnam Journal of Mathematics*, 42(1):83–113, 2014.
- [JRT04] P. Jaillet, E.I. Ronn, and S. Tompaidis. Valuation of commodity-based swing options. *Management Science*, 50:909–921, 2004.
- [Kep04] J. Keppo. Pricing of electricity swing options. *The Journal of Derivatives*, 11(3):26–43, 2004.
- [KH92] M. Kleiber and T.D. Hien. *The Stochastic Finite Element Method*. John Wiley & Sons Ltd, 1992.
- [LK10] O. Le Maître and O.M. Knio. *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*. Springer Science & Business Media, 2010.
- [Loh96] W.L. Loh. On latin hypercube sampling. *Annals of Statistics*, 24(5):2058–2080, 1996.
- [LYL03] W. Liu, Y. Yang, and G. Lu. Viscosity solutions of fully nonlinear parabolic systems. *Journal of Mathematical Analysis and Applications*, 281(1):362–381, 2003.
- [MR98] R. Mikulevicius and B. Rozovskii. Linear parabolic stochastic PDE and Wiener chaos. *SIAM Journal on Mathematical Analysis*, 29(2):452–480, 1998.
- [MR05] R. Mikulevicius and B. Rozovskii. Global L2-solutions of stochastic Navier-Stokes equations. *The Annals of Probability*, 33(1):137–176, 2005.
- [MT02] M. Manoliu and S. Tompaidis. Energy futures prices: term structure models with Kalman filter estimation. *Applied Mathematical Finance*, 9(1):21–43, 2002.
- [Nie92] H. Niederreiter. *Random number generation and quasi-Monte-Carlo methods*, volume 63 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, Philadelphia, PA, 1992.
- [Rao99] B.L.S. Prakasa Rao. *Statistical inference for diffusion type processes, Kendall’s Library of statistics (Vol. 8)*. London/New York: Edward Arnold. Oxford University Press, 1999.

- [Sch97] E. Schwartz. The stochastic behavior of commodity prices: implications for valuation and hedging. *J. Finance*, 52(3):923–973, 1997.
- [SS00] E. Schwartz and J.E. Smith. Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911, 2000.
- [Tou13] A. Tourin. An introduction to finite difference methods for PDEs in finance. In *Optimal Stochastic Target problems and Backward SDE, Fields Institute Monographs*, ed. Nizar Touzi. Springer, 2013.
- [TSGU13] A.L. Teckentrup, R. Scheichl, M.B. Giles, and E. Ullmann. Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numerische Mathematik*, 125(3):569–600, 2013.
- [War16] X. Warin. Some non monotone schemes for time dependent Hamilton-Jacobi-Bellman equations in stochastic control. *J. Sci. Comp.*, 66(3):1122–1147, 2016.
- [Xiu09] D. Xiu. Fast numerical methods for stochastic computations: a review. *Communications in Computational Physics*, 5(2-4):242–272, 2009.