ARTICLE IN PRESS
JID: EOR [m5G;July 28, 2018;15:39]

European Journal of Operational Research 000 (2018) 1–14

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

A general framework for pricing Asian options under stochastic volatility on parallel architectures

Stefania Corsaro[a], Ioannis Kyriakou[b], Daniele Marazzina[c,*], Zelda Marino[a]

[a] Dipartimento di Studi Aziendali e Quantitativi, Università di Napoli Parthenope, via Generale Parisi 13, I-80132, Italy
[b] Cass Business School, City, University of London, 106 Bunhill Row, London EC1Y 8TZ, UK
[c] Dipartimento di Matematica, Politecnico di Milano, via Bonardi 9, Milano I-20133, Italy

ARTICLE INFO

Article history:
Received 30 August 2017
Accepted 7 July 2018
Available online xxx

2010 MSC:
91G20
91G60
65Y05

Keywords:
Finance
Parallel computing
Option pricing
Asian option
Stochastic volatility

ABSTRACT

In this paper, we present a transform-based algorithm for pricing discretely monitored arithmetic Asian options with remarkable accuracy in a general stochastic volatility framework, including affine models and time-changed Lévy processes. The accuracy is justified both theoretically and experimentally. In addition, to speed up the valuation process, we employ high-performance computing technologies. More specifically, we develop a parallel option pricing system that can be easily reproduced on parallel computers, also realized as a cluster of personal computers. Numerical results showing the accuracy, speed and efficiency of the procedure are reported in the paper.

1. Introduction

Transform methods have driven much innovation in financial engineering (see examples of application for different purposes in Carr & Madan, 1999, Albanese, Jackson, & Wiberg, 2004, Feng & Lin, 2013, Cai, Kou, & Liu, 2014, Sesana, Marazzina, & Fusai, 2014, Fusai, Germano, & Marazzina, 2016, Cui, Kirkby, & Nguyen, 2017) and enjoyed wide use in operations research, particularly in queueing theory (see, e.g., Abate & Whitt, 1992, 1995, Abate, Choudhury, & Whitt, 2013, Drekic & Stanford, 2001) with applications as diverse as health care and computer systems' performance evaluation, but also in probability, insurance and radio engineering. It is the scope of the current research to advance the applicability of this indispensable tool by making use of computer technology.

In this paper, we revisit the long-standing problem of valuing non-linear derivatives contingent on the arithmetic average of the underlying asset prices with general monitoring frequency over a prespecified time period. Asian options are very popular among derivatives traders and risk managers, mainly due to the averaging's smoothing of possible market manipulations near the expiry date. Averaging also provides volatility reduction and better cash-flow matching to firms facing streams of cash flows. For this, they appear in currency, energy, metal, agricultural and freight markets and, unsurprisingly, represent a large fraction of the options traded in these markets. Nevertheless, arithmetic averages see wider application in many fields of finance, such as in project valuation (see Zahra & Reza, 2012), optimal capacity planning under average demand uncertainty (see Driouchi, Bennett, & Battisti, 2006), stock-swap merger proposals (see Officer, 2006), technical analysis and algorithmic trading (see Zhu & Zhou, 2009, Kim, 2007).

There is a large body of literature on Asian options; examples include Sesana et al. (2014), Cai, Song, and Kou (2015), Cui, Lee, and Liu (2018), and Černý and Kyriakou (2011) under general Lévy and/or local volatility models. Chang and McAleer (2015) study Garch-like models, which are closely linked with econometric analysis, and their relation to option pricing; examples of option applications include Majewski, Bormetti, and Corsi (2015), but also Mercuri (2011) on Asian options. For a more detailed literature survey, we refer to Fusai and Kyriakou (2016).

In this paper, we consider a general stochastic volatility model framework with or without asset price jumps. Stochastic volatility is a salient and well-documented feature of financial assets. Empirical results of Bakshi, Cao, and Chen (1997) suggest that stochastic

* Corresponding author.
E-mail addresses: stefania.corsaro@uniparthenope.it (S. Corsaro), ioannis.kyriakou@city.ac.uk (I. Kyriakou), daniele.marazzina@polimi.it (D. Marazzina), zelda.marino@uniparthenope.it (Z. Marino).

https://doi.org/10.1016/j.ejor.2018.07.017

Please cite this article as: S. Corsaro et al., A general framework for pricing Asian options under stochastic volatility on parallel architectures, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.07.017

volatility is a primary factor driving option prices. Diffusion-based volatility models account for dependence in increments and long-term smiles and skews, but cannot give rise to realistic short-term implied volatility patterns; this shortcoming can be overcome by introducing jumps in the returns (see Cont & Tankov, 2004). Under such models, analytical price solutions for Asian options are approximate and scarce, including Yamazaki (2014), Zeng and Kwok (2016), and Fusai and Kyriakou (2016). Instead, we resort to a recursive transform approach, which is free from restricted closeness to the true option price but rather enjoys superior efficiency that has been evidenced, for example, by Carverhill and Clewlow (1990) and Černý and Kyriakou (2011) under simpler model constructions.

We present a parallel system aiming to enhance the computational tractability of the valuation problem, an important aspect of concern of transform techniques in operations research (e.g., see Drenovak, Ranković, Ivanović, Urošević, & Jelic, 2017). By applying the parallel system to the numerical pricing procedure, we show that it leads to fast computation which further allows the users to exploit its smooth convergence, thereby producing price results at any desired accuracy level. We also conduct a theoretical analysis of the errors due to truncation and discretization involved in the transforms.

Our parallel system is inexpensive as it comprises a Non-Uniform Memory Access machine equipped by open-source software at all levels, including the operating system, the compiler, the message-passing system and the routines for main computational kernels arising in the pricing procedure. The hardware architecture is representative of many recent high-performance computing (HPC) ones. We recognize in them a mixed-memory (distributed and shared) architecture. The processors are typically grouped by their physical location on a node, a multi-core Central Processing Unit (CPU) package, and the nodes are coupled via high-speed interconnects. We point out that modern personal computers (PCs) have multi-core processors equipped with local cache memory; thus, mixed-memory models adapt to them, too. The option pricing system can then be easily ported on parallel computers, also realized as a cluster of PCs.

The proposed parallel procedure enjoys flexibility as it is built on sophisticated asset price model dynamics reflecting the market realism, while remaining computationally tractable, hence advancing the universality of transform methods in this direction. In particular, the method relies on knowledge of the characteristic function of the (log) price conditional on the current and future variance states, which we derive for different affine stochastic volatility models including, among others, time changed Lévy processes (Carr, Geman, Madan, & Yor, 2003) as well as the classical Heston (1993) and Bates (1996) models. We bypass increases in computational complexity and memory allocation due to increasing dimensionality by employing HPC technologies. HPC is recognized as mandatory for the effective solution of many financial problems when complex simulations must be performed in a suitable turnaround time; see, for example, Östermark (2017) where a recursive portfolio decision system is described. In this paper, we employ *domain decomposition*, according to which the computational domain is split among the involved processes. Then, each process works on a local sub-domain, communicating with the others when required. Clearly, a good strategy should require limited communication, so as to minimize the impact of communication overhead on performance. We show that our strategy leads to an inherently parallel procedure: the grid is split in such a way that processes work concurrently on a local sub-grid most of the time. In particular, communication is not required in the transform computation, which is crucial for the sake of parallel efficiency. We show that efficiency values close to the ideal ones are feasible.

The remainder of the paper is organized as follows. Section 2 presents our financial modeling framework and main preliminary results. Section 3 develops the backward-recursive integral pricing scheme, whereas Section 4 discusses its numerical implementation and introduces the parallel system. Section 5 provides a theoretical error analysis of the scheme. Section 6 presents some numerical experiments, and Section 7 concludes the paper. Proofs and additional results are collected in appendices.

2. Stochastic volatility framework

Let $(\Omega, \mathcal{F}, \hat{P})$ be a complete probability space upon which all stochastic processes are defined. We denote by $\hat{P}$ the risk neutral probability measure.

When the risk neutral dynamics of the log-price is given by a process with independent increments, the implied volatility surface follows a deterministic evolution (see Cont & Tankov, 2004); introducing stochastic volatility can tackle this difficulty. Such models are popular among practitioners and academics as they are able to account for volatility clustering and dependence in increments, and give rise to realistic implied volatility patterns (short-term and/or long-term skews). In this work, we study models from the affine class, such as Heston (1993), Bates (1996) and Carr et al. (2003), although non-affine models (e.g., see Hagan, Kumar, Lesniewski, & Woodward, 2002) represent an interesting direction for future research.

2.1. Diffusion-based stochastic volatility

In the Heston model, the stochastic variance V follows the CIR (named after Cox, Ingersoll, & Ross, 1985; see also Feller, 1951) mean reverting, square root diffusion model

$$dV_t = \alpha(\beta - V_t)dt + \gamma\sqrt{V_t}dW_t^V, \ V_0 = \upsilon_0, \tag{1}$$

where $(W_t^V)_{t\geq 0}$ is a standard Brownian motion. The parameters α, β, γ, υ_0 are positive constants. Process (1) is continuous and non-negative. Its exact behavior near zero depends on the parameters' values: if $\gamma^2 \leq 2\alpha\beta$, the process never touches zero. The log-price under the risk neutral measure $\hat{P}$ is given by

$$dX_t = \left(r - \frac{1}{2}V_t\right)dt + \sqrt{V_t}\left(\rho dW_t^V + \sqrt{1-\rho^2}dW_t^X\right), \tag{2}$$

where r is the continuously compounded risk free interest rate and $(W_t^X)_{t\geq 0}$ a standard Brownian motion. W^V and W^X are independent. Finally, ρ is the correlation coefficient between processes V and X: $\rho < 0$ leads to the so-called leverage effect of large downward price moves being associated with upward volatility moves, and is offered as an explanation for implied volatility skews (e.g., see Cont & Tankov, 2004).

2.2. Stochastic volatility model with price jumps

Diffusion-based stochastic volatility models cannot generate sufficient asymmetry in the short-term returns, hence are unable to match the empirical short-term skews. The jump diffusion stochastic volatility model of Bates (1996) addresses this shortcoming by superimposing price jumps (see Cont & Tankov, 2004). In addition, the leverage effect and the long-term skews are achieved using correlated Brownian motions in the price and variance processes. More specifically, in this model, the risk neutral log-price is given by

$$\begin{aligned} dX_t = &\left(r - \lambda(e^{k(1)} - 1) - \frac{1}{2}V_t\right)dt \\ &+ \sqrt{V_t}\left(\rho dW_t^V + \sqrt{1-\rho^2}dW_t^X\right) + dL_t, \end{aligned} \tag{3}$$