MONTE CARLO BOUNDS FOR GAME OPTIONS INCLUDING CONVERTIBLE BONDS

CHRISTOPHER BEVERIDGE AND MARK JOSHI

Abstract. We introduce two new methods to calculate bounds for zero-sum game options using Monte Carlo simulation. These extend and generalise the duality results of Haugh–Kogan/Rogers and Jamshidian to the case where both parties of a contract have Bermudan optionality. It is shown that the Andersen–Broadie method can still be used as a generic way to obtain bounds in the extended framework, and we apply the new results to the pricing of convertible bonds by simulation.

1. Introduction

Game options represent an important class of options. They have the distinguishing feature that both the holder and the issuer have early exercise rights, and can therefore be characterized as a two-player extensive form game. It is this complex feature that makes the pricing of game options particularly difficult, and as a consequence, the techniques used to handle game options are lagging behind those used to handle simpler options. In this paper, we address the problem of pricing zero-sum game options in which all cash-flows occur solely between the two counterparties. In particular, we show how Monte Carlo simulation can be used to obtain unbiased estimates of upper and lower bounds from a given pair of exercise strategies.

Arguably the most important examples of such game options are convertible bonds. These are complex hybrid derivatives, subject to equity, interest rate and credit risk. A typical convertible bond consists of a coupon-paying bond, which the holder can convert to shares at any time. In addition, the holder can usually put the bond back to the issuer on a discrete set of dates, and the issuer can call the bond early, possibly subject to some form of call protection. See Grimwood and Hodges (2002) for a good summary of the precise structure of typical convertible bonds.

We restrict our attention to games where exercise can occur precisely once, thus the game terminates as soon as one of the players decides. Importantly, we assume that one party has priority when exercising and focus on the case of complete and perfect information; that is, each player has all the information available in the market and knows the other player's optimal course of action. This simplifies the analysis, and in particular the underlying exercise strategies, whilst providing a good approximation to reality. In particular, it is sufficiently wide to encompass convertible bonds which often carry the provision that the holder can convert even if the bond is to be called.

To be able to handle the callability, much of the early work on pricing convertible bonds focused on the use of deterministic backwards methods, including binomial trees and PDE methods; see, for example, Tsiveriotis and Fernandes (1998), Takahashi et al. (2001) and Andersen and Buffum (2002). However, since these methods suffer from the so-called curse of dimensionality, they become relatively inefficient after we have a small number of factors. As such, simplified models are usually required, neglecting certain risk factors. In addition, these backwards methods can struggle to handle the path-dependent call protection common in practice, to the point where they are no longer feasible for particularly nasty forms of call-protection; see Crépey and Rahal (2009).

Riding on the back of recent advances in the Monte Carlo pricing of early-exercisable derivatives where just one party has Bermudan optionality, Monte Carlo algorithms for the pricing of convertible bonds have been introduced by Lvov et al. (2004a) and Crépey and Rahal (2009) to overcome these shortfalls. In particular, since the order of convergence for Monte Carlo simulation is independent of the dimension, and Monte Carlo simulation can easily handle path-dependence, combined with the ability to handle the exercise decisions it is the natural tool for such problems.

However, the Monte Carlo methods for handling game options are still in their infancy relative to those used where only one party has the callability. The algorithms of Lvov et al. (2004a) and Crépey and Rahal (2009) extend the least-squares method introduced by Longstaff and Schwartz (2001) and Carrière (1996). Whilst Chassagneux and Crépey (2009) proved results on the rate of convergence, there is still a need for methodologies that establish unbiased estimates of upper and lower bounds on the price. If these are tight then one can be sure that the price is accurate.

This is in marked contrast to the case where only one party can exercise, since then any two-pass methodology, such as least-squares, that first estimates an exercise strategy and then uses a second independent simulation to estimate the price yields an unbiased estimate of a lower bound. It is also possible to calculate upper bounds using the duality approaches of Rogers (2002), Haugh and Kogan (2004), and Jamshidian (2004); see Chen and Glasserman (2007) for a comparison of the duality approaches. In addition, for game options, since the least-squares method no longer gives a lower bound, it is hard even to assess what direction the error is in and whether a change to the development of the exercise strategy is an improvement.

In this paper, we introduce two new methods that allow the calculation of both lower and upper bounds for game options using Monte Carlo simulation. To do this, we extend the duality approaches used for single early-exercisable derivatives. One method is based on the additive approach of Davis and Karatzas (1994), Rogers (2002) and Haugh and Kogan (2004), while the other is based on the multiplicative approach of Jamshidian (2004). In the single case, Andersen and Broadie (2004) presented a generic methodology using sub-Monte carlo simulations to turn exercise strategies into unbiased estimates of upper bounds. Here we extend their ideas and present a methodology for turning pairs of exercise strategies for the game option into unbiased estimates of upper and lower bounds. As in the single-exercisable case, this method has the additional virtue that the accuracy of the exercise strategies directly translates into the tightness of the bounds obtained. In other words, if the bounds obtained are far apart then the exercise strategies are poor.

This is an important step forward, as it provides an easy way to assess the accuracy of the prices obtained via the methods introduced by Lvov et al. (2004a) and Crépey and Rahal (2009), or indeed by any methodology that estimates exercise strategies.

We structure the paper as follows. In Section 2, we introduce notation. Section 3 contains the main results of the paper, giving results which can be used to calculate bounds for game options.

Section 4 discusses the practical implementation of these results, and Section 5 contains numerical results.

We are grateful to two anonymous referees for their detailed comments on earlier versions of this paper.

2. Notation

Assume we have a fixed time horizon, T, and a filtered probability space $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{Q})$ satisfying the usual properties, where $\mathbb{Q}$ denotes an equivalent martingale measure. In addition, assume that $\mathbb{Q}$ is the measure associated with using a general numeraire whose value is denoted by $N(t)$. Throughout, any reference to deflated values will mean that the values are measured in units of this numeraire asset.

Consider a game option with two participants, A and B, who both have Bermudan optionality. We focus on zero-sum game options where all cash-flows occur between A and B. In addition, we restrict our attention to games where exercise can occur precisely once and one party has priority over the other when exercising. That is, if A and B both decide to exercise at a particular time, we assume one party is allowed to exercise before the other, and that the game terminates at that moment, before the second player can exercise. We assume that the option generates a sequence of cash-flows until either player exercises, or default occurs. Upon exercise or default, a rebate is exchanged between A and B, where the size of the rebate depends on who exercises or whether default occurs. As mentioned above, to simplify the analysis we focus on the case of complete and perfect information; that is, each player has all the information available in the market and knows the other player's optimal course of action. For concreteness, we shall assume that B has priority over A. Note that this implies that B's decision as to whether to exercise at a given time will vary according to A's decision. In particular, if A has exercised then B exercises if and only if his exercise value is greater than it is on A's exercise, whereas if A has not then B exercises if and only if his exercise value is greater than his continuation value. It may also be to A's advantage to exercise even if he knows that B will counter-exercise immediately, if the pay-off from B's exercise is smaller than B's continuation value in the event of neither party exercising.

A consequence of this is that if B has priority over A and A exercises, then A knows that B will immediately exercise, using his priority, if and only if the pay-off is to his advantage. Thus by redefining the pay-off on A's exercise to be the minimum (from A's viewpoint) of the two pay-offs, we can replace the game by an equivalent one in which the pay-off does not depend on the identity of the exerciser. This simplifies the analysis and we therefore derive our theoretical results in this context.

Let τ_A denote the exercise time of A and let τ_B denote that of B, where both are assumed to be stopping times. Also, we allow for the possibility of default by assuming that τ_d denotes an exogenously given stopping time representing the time of default. We emphasize that default cannot be influenced by either A or B in this model.

We redefine the exercise value at maturity to agree with the non-exercise value in cases where no exercise would occur. This allows us to force exercise at maturity without changing the product's value. The time of exercise then becomes a bounded stopping time which is convenient for the proofs.

Let $\Xi^A = \{T_1, \ldots, T_m\}$ denote the set of exercise times for A, where

$$0 = T_0 < T_1 < \ldots < T_m = T.$$

We allow each T_i to be a stopping time to allow for exercise protection, such as no-call periods and more complicated path-dependent forms. In addition, let Γ_t^A denote the set of possible exercise strategies for A taking values in $[t, T] \cap \Xi^A$, and Ω_t^A denote the set of random times taking values in $[t, T] \cap \Xi^A$.

For simplicity, throughout we consider cash-flows and therefore prices from A's point of view. In particular, let D_j denote the deflated cash-flows occurring in $(T_{j-1}, T_j]$, ceased at the time of default or exercise by B. Since by definition D_j gives the cash-flows occuring between A's exercise times, it can not be influenced by A.

Define $f_t(\tau_A, \tau_B)$ to be the deflated cash-flows received by A on $[t, T]$. For technical reasons, we assume that

$$\sup_{s \in \Xi^A} \sup_{t \in \Xi^B} |f_0(s, t)| \in L^1. \tag{2.1}$$

Finally, let $V(t)$ denote the deflated value of the convertible bond at time t, so

$$\begin{aligned} V(t) &= \sup_{\tau_A \in \Gamma_t^A} \inf_{\tau_B \in \Gamma_t^B} \mathbb{E}\left[f_t(\tau_A, \tau_B) | \mathcal{F}_t\right], \\ &= \inf_{\tau_B \in \Gamma_t^B} \sup_{\tau_A \in \Gamma_t^A} \mathbb{E}\left[f_t(\tau_A, \tau_B) | \mathcal{F}_t\right], \end{aligned} \tag{2.2}$$

see, for example, Bielecki et al. (2008) and the references therein.

We focus on the pricing of game options, and have therefore restricted our attention to the case where the price is given by the expectation of deflated cash-flows in an equivalent martingale measure. However, for the additive method all that really matters is that the value is given by (2.2), and so if this and the other requirements described above hold, then the martingale measure is not strictly necessary.

Although we will focus on applying our results to convertible bonds, we have kept the treatment general to emphasize that our results can be applied to a variety of products in a range of models. For example, cancellable swaps, where cancellation is possible by either side but involves paying a fee, priced in the LIBOR market model can easily fit into the framework presented. Since both parties have Bermudan optionality, such a cancellable swap is a zero-sum game option, and all that we have to do is define the payoff function, f, appropriately.

3. Bounds

In this section, we derive results which allow the calculation of bounds for game options. These results generalise the standard duality approaches of Rogers (2002), Haugh and Kogan (2004) and Jamshidian (2004); by removing the Bermudan optionality of one party, the original results are obtained.

It is common when pricing derivatives contracts to study a continuous time model which is then discretized for pricing purposes in a computer. For example, when pricing an Asian option in the Black–Scholes model, the pricing model is justified using continuous time arguments and these reduce the pricing to the computation of a finite-dimensional integral. The integral is then evaluated by using a discretization of the model in a computer. It is therefore a matter of judgement and taste at what point the discretization occurs. Here we specify continuous-time models, but justify our upper bounds in the discretized setting following Jamshidian (2004). This saves us from having to rephrase our results for the discrete setting when discussing how to implement the model using Monte Carlo simulation. However, we note that our results could be extended to continuous time through a suitable modification of the arguments used by Rogers (2002).

3.1. **Additive Bounds.** Our first approach to finding bounds for game options extends the additive approach of Rogers (2002) and Haugh and Kogan (2004). See also Davis and Karatzas (1994) for initial results, Joshi (2006, 2007) for generalisations and improvements, and Glasserman (2004) for a good overview.

Since both parties have callability and we are modeling zero-sum games, the problem is symmetric. In particular, any arguments applied to the valuation of game options from A's perspective can be applied without modification by considering the same option from the perspective of B. As such, if we can obtain upper bounds from A's perspective, then we can also get upper bounds from B's perspective using the same arguments. Since an upper bound from B's perspective is actually a lower bound from A's perspective, it is enough to show how to obtain one-sided bounds from A's point of view. Therefore, we do not overburden the reader with unnecessary exposition, restricting ourselves to the case of obtaining upper bounds from A's perspective.

First we show how to actually obtain bounds. Assume that we have a process, M, with initial value zero that is a martingale in our pricing measure, and let $\hat{\tau}_B$ denote any fixed exercise strategy for B. From (2.2),

$$\begin{aligned} V(0) &= \inf_{\tau_B \in \Gamma_0^B} \sup_{\tau_A \in \Gamma_0^A} \mathbb{E}\left[f_0(\tau_A, \tau_B)\right], \\ &\leq \sup_{\tau_A \in \Gamma_0^A} \mathbb{E}\left[f_0(\tau_A, \hat{\tau}_B)\right], \end{aligned}$$

where the inequality follows since B is using a fixed exercise strategy which may be sub-optimal. Now, since exercise times are bounded and $M(0) = 0$, the optional stopping theorem gives

$$V(0) \leq \sup_{\tau_A \in \Gamma_0^A} \mathbb{E}\left[f_0(\tau_A, \hat{\tau}_B) - M(\tau_A)\right],$$

and taking the supremum over the set of random times rather than stopping times gives,

$$\begin{aligned} V(0) &\leq \sup_{\tau_A \in \Omega_0^A} \mathbb{E}\left[f_0(\tau_A, \hat{\tau}_B) - M(\tau_A)\right], \\ &= \mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \hat{\tau}_B) - M(t_k)\right)\right]. \end{aligned} \tag{3.1}$$

This expression is what we are looking for; given an exercise strategy for B, we have an expectation which we can evaluate to obtain an upper bound for the game option. While this is the actual expression we use to calculate bounds, we now show that it is indeed possible to obtain equality when M is chosen optimally.

To do this, we fix τ_B and let

$$\hat{V}(T_j, \tau_B) := \sup_{\tau_A \in \Gamma^A_{T_j}} \mathbb{E}\left[f_{T_j}(\tau_A, \tau_B) | \mathcal{F}_{T_j} \right].$$

Consider the martingale part of the Doob decomposition of $\hat{V}(T_j, \tau_B) + \sum_{k=1}^{j} D_k$, which we denote by $\bar{M}(T_k, \tau_B)$. Since (2.1) ensures that $\hat{V}(T_j, \tau_B) + \sum_{k=1}^{j} D_k \in L^1$ for all j and it is adapted, it has a Doob decomposition with $\bar{M}(T_k, \tau_B)$ given by

$$\bar{M}(T_k, \tau_B) := \Delta_1 + \ldots + \Delta_k, \tag{3.2}$$

where,

$$\Delta_k := \hat{V}(T_k, \tau_B) + D_k - \mathbb{E}\left[\hat{V}(T_k, \tau_B) + D_k | \mathcal{F}_{T_{k-1}} \right]. \tag{3.3}$$

Using this martingale, we get the following.

Theorem 1.

$$V(0) = \inf_{\tau_B \in \Gamma^B_0} \mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - \bar{M}(t_k, \tau_B) \right) \right]. \tag{3.4}$$

Proof. We use induction to show a.s. that

$$\hat{V}(T_k, \tau_B) = \max\left(f_{T_k}(T_k, \tau_B), f_{T_k}(T_{k+1}, \tau_B) - \Delta_{k+1}, \ldots, f_{T_k}(T_m, \tau_B) - \Delta_{k+1} - \ldots - \Delta_m \right). \tag{3.5}$$

This holds at $k = m$ because, by definition of $f_{T_m}(T_m, \tau_B)$, $\hat{V}(T_m, \tau_B) = f_{T_m}(T_m, \tau_B)$. Now assume it holds at k. Using the definition of $\hat{V}(T_{k-1}, \tau_B)$ first and then using (3.3) to write in terms of Δ_k,

$$\begin{aligned} \hat{V}(T_{k-1}, \tau_B) &= \max\left(f_{T_{k-1}}(T_{k-1}, \tau_B), \mathbb{E}\left[\hat{V}(T_k, \tau_B) + D_k | \mathcal{F}_{T_{k-1}} \right] \right), \\ &= \max\left(f_{T_{k-1}}(T_{k-1}, \tau_B), \hat{V}(T_k, \tau_B) + D_k - \Delta_k \right). \end{aligned}$$

By the inductive hypothesis, we can substitute (3.5) for $\hat{V}(T_k, \tau_B)$ and absorb D_k into the payoff unction f, giving

$$\hat{V}(T_{k-1}, \tau_B) = \max\left(f_{T_{k-1}}(T_{k-1}, \tau_B), f_{T_{k-1}}(T_k, \tau_B) - \Delta_k, \ldots, f_{T_{k-1}}(T_m, \tau_B) - \Delta_k - \ldots - \Delta_m \right).$$

Therefore, (3.5) holds for $0 \leq k \leq m$ by induction.

In particular, by setting $k = 0$ in (3.5) and substituting in the definition of $\bar{M}(T_k, \tau_B)$,

$$\hat{V}(0, \tau_B) = \max_{T_k \in \Xi^A} \left(f_0(T_k, \tau_B) - \bar{M}(T_k, \tau_B) \right). \tag{3.6}$$

Since by definition,

$$V(0) = \inf_{\tau_B \in \Gamma^B_0} \hat{V}(0, \tau_B),$$

we get (3.4). $\square$

Remark 1. *The above proof actually shows that*

$$V(0) = \inf_{\tau_B \in \Gamma_0^B} \max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - \bar{M}(t_k, \tau_B) \right), \text{ a.s.},$$

that is, the result of Theorem 1 holds without expectation because the optimal martingale leads to zero variance.

So not only can we obtain bounds from (3.4), but if τ_B and M are chosen optimally, then it is indeed possible to obtain equality with the true price. We present an alternative method for calculating bounds in the next section.

3.2. **Multiplicative Bounds.** Rather than use the additive approach to obtain bounds, we can use the multiplicative approach of Jamshidian (2004). As in the case of standard early exercisable derivatives, the multiplicative approach requires that the value of the underlying product is strictly positive. As such, we employ the additional assumption that a.s.

$$f_t(\tau_A, \tau_B) > 0$$

for all t and any τ_A and τ_B. As a consequence, we cannot apply the symmetry argument we used with the additive approach to restrict our attention to obtaining upper bounds. This is because a positive payoff function from A's point of view will be negative from the perspective of B. However, since lower bounds can be easily obtained with the multiplicative approach in almost the same way as upper bounds, we again restrict our attention to obtaining upper bounds from A's point of view.

Suppose that we have a general numeraire whose value is given by $G(t)$ and let $\hat{\tau}_B$ denote any fixed exercise strategy for B. Changing measure gives,

$$\begin{aligned} V(0)N(0) &= N(0) \inf_{\tau_B \in \Gamma_0^B} \sup_{\tau_A \in \Gamma_0^A} \mathbb{E}\left[f_0(\tau_A, \tau_B) \right], \\ &= G(0) \inf_{\tau_B \in \Gamma_0^B} \sup_{\tau_A \in \Gamma_0^A} \mathbb{E}_G \left[f_0(\tau_A, \tau_B) \frac{N(\min(\tau_A, \tau_B, \tau_d))}{G(\min(\tau_A, \tau_B, \tau_d))} \right], \end{aligned}$$

where $\mathbb{E}_G[.]$ denotes that the expectation is taken in the measure associated with using G as numeraire. Taking the supremum over the set of random times rather than stopping times gives,

$$\begin{aligned} V(0)N(0) &\leq G(0) \inf_{\tau_B \in \Gamma_0^B} \sup_{\tau_A \in \Omega_0^A} \mathbb{E}_G \left[f_0(\tau_A, \tau_B) \frac{N(\min(\tau_A, \tau_B, \tau_d))}{G(\min(\tau_A, \tau_B, \tau_d))} \right], \\ &= G(0) \inf_{\tau_B \in \Gamma_0^B} \mathbb{E}_G \left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) \frac{N(\min(t_k, \tau_B, \tau_d))}{G(\min(t_k, \tau_B, \tau_d))} \right) \right]. \qquad (3.7) \end{aligned}$$

We can interpret the final expression as the price of a product paying one unit of G at time t_n. We can therefore price it with a different numeraire. We therefore return to $\mathbb{Q}$, where we have tractable dynamics for our state variables, and obtain our final result,

$$\begin{aligned} V(0)N(0) &\leq N(0) \inf_{\tau_B \in \Gamma_0^B} \mathbb{E} \left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) \frac{N(\min(t_k, \tau_B, \tau_d))}{G(\min(t_k, \tau_B, \tau_d))} \right) \frac{G(t_n)}{N(t_n)} \right], \qquad (3.8) \\ &\leq N(0) \mathbb{E} \left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \hat{\tau}_B) \frac{N(\min(t_k, \hat{\tau}_B, \tau_d))}{G(\min(t_k, \hat{\tau}_B, \tau_d))} \right) \frac{G(t_n)}{N(t_n)} \right]. \qquad (3.9) \end{aligned}$$

While (3.9) allows us to calculate bounds, using similar arguments to those of Joshi (2007), we show that it is actually possible to obtain equality by choosing an optimal numeraire and exercise strategy for B. Notation from Section 3.1 is used.

Theorem 2. *Define* $\tilde{M}(t,\tau_B) := \left(\bar{M}(t,\tau_B) + \hat{V}(0,\tau_B)\right)$. *Provided* $f_t(\tau_A,\tau_B) > 0$ *for all* t, τ_A *and* τ_B *a.s., we have*

$$V(0) = \inf_{\tau_B \in \Gamma_0^B} \mathbb{E}\left[\max_{t_k \in \Xi^A} \left(\frac{f_0(t_k,\tau_B)}{\tilde{M}(t_k,\tau_B)}\right) \tilde{M}(t_n,\tau_B)\right]. \tag{3.10}$$

Proof. Fix τ_B and consider taking as numeraire the optimal hedge portfolio from Section 3.1, with the slight modification that the hedge portfolio starts with the underlying product and therefore has initial value $N(0)\hat{V}(0,\tau_B)$. In particular, take

$$G^*(t,\tau_B) := N(t)\left(\bar{M}(t,\tau_B) + \hat{V}(0,\tau_B)\right) = N(t)\tilde{M}(t,\tau_B). \tag{3.11}$$

Since $G^*(t,\tau_B)$ denotes the value of a self-financing portfolio (see, for example, Joshi (2006)) which is positive due to (3.6) since $f_t(\tau_A,\tau_B) > 0$ a.s., by the results of Geman et al. (1995) it is a valid numeraire.

Now, from (3.4) and (3.6),

$$\tilde{M}(T_k,\tau_B) \geq f_0(T_k,\tau_B),$$

for all $T_k \in \Xi_0^A$ a.s. Therefore

$$\frac{f_0(T_k,\tau_B)}{\tilde{M}(T_k,\tau_B)} \leq 1,$$

which implies that

$$\mathbb{E}_{G^*}\left[\max_{T_k \in \Xi^A}\left(f_0(T_k,\tau_B)\frac{N(T_k)}{G^*(T_k,\tau_B)}\right)\right] \leq 1,$$

since $\frac{N(t)}{G^*(t,\tau_B)} = \frac{1}{\tilde{M}(t,\tau_B)}$ from (3.11). Finally,

$$\hat{V}(0,\tau_B)N(0) \geq G^*(0,\tau_B)\mathbb{E}_{G^*}\left[\max_{T_k \in \Xi^A}\left(f_0(T_k,\tau_B)\frac{N(T_k)}{G^*(T_k,\tau_B)}\right)\right],$$

since $\hat{V}(0,\tau_B)N(0) = G^*(0,\tau_B) > 0$ from (3.11). Combining this with (3.7) yields

$$\hat{V}(0,\tau_B)N(0) = G^*(0,\tau_B)\mathbb{E}_{G^*}\left[\max_{T_k \in \Xi^A}\left(f_0(T_k,\tau_B)\frac{N(T_k)}{G^*(T_k,\tau_B)}\right)\right].$$

We can then perform another measure change, returning to $\mathbb{Q}$, where we have tractable dynamics for the state variables of our model. In particular, writing G^* in terms of N and $\tilde{M}$,

$$\hat{V}(0,\tau_B) = \mathbb{E}\left[\max_{T_k \in \Xi^A}\left(\frac{f_0(T_k,\tau_B)}{\tilde{M}(T_k,\tau_B)}\right)\tilde{M}(T_m,\tau_B)\right],$$

and since, by definition,

$$V(0) = \inf_{\tau_B \in \Gamma_0^B} \hat{V}(0,\tau_B),$$

we get (3.10). □

Therefore, we have another way of obtaining bounds. In particular, from (3.8) and (3.9) we see that any portfolio other than the optimal one will lead to an upward bias, yet we can obtain equality as in (3.10). The restriction that $\tilde{M}(t, \tau_i)$ remain positive will usually be satisfied in practice for convertible bonds. This follows since from (3.6), $\tilde{M}(t, \tau_B)$ is never less than the value of the derivative, which is in turn strictly positive if we assume non-zero recovery upon default.

4. Practical Implementation

Although we have obtained results that allow for the calculation of bounds, to actually use these results we need to develop accurate approximations to the optimal exercise strategies for A and B and for the optimal martingales. Developing approximate exercise strategies has been taken care of by Lvov et al. (2004a) and Crépey and Rahal (2009), who extended the least-squares method to apply to game options. In this section, we discuss how to improve the accuracy of the least-squares approximate exercise strategies for convertible bonds, and how to develop accurate approximations to the optimal martingales. For completeness, we start with a detailed description of the least-squares method in the context of game options.

4.1. **The Least-Squares Method.** Lvov et al. (2004a) and Crépey and Rahal (2009) extended the least-squares method of Longstaff and Schwartz (2001) and Carrière (1996) to the pricing of game options. Their approach can be used to develop approximate exercise strategies for A and B, which can then be used to approximate the price of game options. Here we describe this method.

To use the least-squares method, we first have to specify a set of basis functions which are used to approximate the continuation value at each exercise time. For convertible bonds, these could be polynomials in the stock price and interest rates for example. Once we have done this, the algorithm for the least-squares method can be broken into three main components: the first pass, the regression component, and the second pass.

For the first pass we run a simulation, storing at each exercise time on each path the deflated exercise values for A and B, the deflated cash-flows occurring between exercise dates, the value of the basis functions and whether default has occurred or not. To account for different forms of exercise protection, if A cannot exercise, set A's exercise value to negative infinity and if B cannot exercise, set B's exercise value to positive infinity. This is precisely the information we require to develop our approximate exercise strategy. To make things concrete, assume that M_1 paths are used.

Once this is done, we move on to the regression component. We start at the final exercise time and work back iteratively, developing our approximate exercise strategy at each exercise time as we go. In particular, without loss of generality, assume that we are at the j^{th} exercise time. Before we start our description, consider the following notation. Let,

- $H^i_{A,j}$ denote A's deflated exercise value for the i^{th} path and j^{th} exercise time, with similar notation for B,
- D^i_j denote the deflated cash-flows received by A between the previous and current exercise times for the i^{th} path,
- $BF^i_{j,k}$ denote the value of the k^{th} basis function on the i^{th} path,
- Z^i_j denote observations of the deflated continuation value from the j^{th} exercise time, where exercise at future dates occurs according to our recently developed approximate strategies.

Assume we have M_{BF} basis functions. Of course it is possible to use different basis functions at each exercise time. However, we ignore this trivial generalisation to ease exposition. In addition, note that at the final exercise date or the exercise date before default occurs, Z_j^i will just be the observed cash-flows occurring after that date, and that, as we move backwards through the exercise times, we can update our pathwise observations of the continuation value with our approximate exercise strategies.

We develop our approximate exercise strategies by first developing an estimate to the continuation value. We approximate the deflated continuation value as,

$$\hat{V}_j = \sum_{k=1}^{M_{BF}} \lambda_{j,k} BF_{j,k},$$

where the coefficients, $\lambda_{j,k}$, are estimated by a least-squares regression. In particular, our basis functions are regressed against our observed continuation values, Z_j^i, for each path where default has not occurred. That is, the $\lambda_{j,k}$ are chosen to minimise

$$\sum_{i=1}^{M_1} \left(Z_j^i - \sum_{k=1}^{M_{BF}} \lambda_{j,k} BF_{j,k}^i \right)^2 Y_j^i \mathrm{I}_{\left\{H_{A,j}^i < \infty\right\} \cup \left\{H_{B,j}^i > -\infty\right\}},$$

where Y_j^i takes the value zero if default has occurred before the j^{th} exercise time on the i^{th} path and one otherwise. Note that such a minimisation problem has a simple solution, which involves solving an $M_{BF} \times M_{BF}$ linear system of equations. As such, the least-squares regressions can be implemented very efficiently.

Given that we can now estimate the continuation value, it is possible to back-out our approximate exercise strategies. Each party takes the course of action which maximises their holdings; that is, they exercise if the estimated continuation value is less than their exercise value. From this exercise strategy, we can update the pathwise observations of the continuation value, to give observations from the next exercise date backwards in time as follows:

$$Z_{j-1}^i = \begin{cases} D_j^i + H_{A,j}^i, & \text{if } H_{A,j}^i \geq \min\left(\hat{V}_j^i, H_{B,j}^i\right), \\ D_j^i + H_{B,j}^i, & \text{if } H_{A,j}^i < \min\left(\hat{V}_j^i, H_{B,j}^i\right) \text{ and } H_{B,j}^i \leq \hat{V}_j^i, \\ D_j^i + Z_j^i, & \text{otherwise,} \end{cases}$$

where we have assumed that A can exercise before B. This is done for each path where default has not yet occurred. Noting that all cash-flows are from A's point of view, the first line in the update rule for Z_j^i follows since A will exercise if either A's exercise value is above the continuation value, or if B exercises (that is, $H_{B,j}^i \leq \hat{V}_j^i$) and A's exercise value is above B's. This is encapsulated in the rule $H_{A,j}^i \geq \min\left(\hat{V}_j^i, H_{B,j}^i\right)$. In addition, B will exercise if A does not exercise and it is optimal for B to exercise, leading to the second line. Finally, if neither A or B exercise, the observed continuation value for path i will just be that for the next time frame plus any intermediate cash-flows occuring between the exercise dates.

We repeat the above procedure until we have finished with t_1. It is then possible to obtain an initial estimate of the option's price using

$$N(0)\frac{1}{M_1}\sum_{i=1}^{M_1} Z_0^i.$$

However, this gives a biased estimate of the price since information ahead of the current exercise time is implicitly used to evaluate the exercise decision, introducing foresight bias; see the paper by Fries (2005).

To remove this bias, it is common to perform the third step of the algorithm, often referred to as the second pass. In particular, an independent pricing simulation is used to estimate

$$N(0)\mathbb{E}\left[f_0\left(\tau_A^{\mathrm{LS}}, \tau_B^{\mathrm{LS}}\right)\right],$$

where τ_A^{LS} and τ_B^{LS} are the least-squares approximate strategies.

4.2. **Improving the Least-Squares Exercise Strategies for Convertible Bonds.** In Lvov (2005) and Lvov et al. (2004b), a basic least-squares procedure, such as the one above, is advocated. However, in Lvov et al. (2004a), an improvement to the least-squares method for convertible bonds is briefly mentioned, which we now describe.

Convertible bonds behave very differently depending on the value of the underlying stock. For very high stock prices, where conversion is likely to occur, a convertible bond is very similar to equity. In contrast, for very low stock prices, convertible bonds behave much more like corporate bonds. As such, in Lvov et al. (2004a) the authors warn that using a global least-squares regression is likely to be rather ineffective.

A simple solution is to partition the stock price's state space, and use separate least-squares regressions for the different regions. This should allow better fits to the continuation value in the different regions relative to using one global regression, and therefore lead to improved exercise strategies. This extends the double regression enhancement introduced by Beveridge and Joshi (2009), and we will refer to it as the multiple regression enhancement. In addition, it should not be very expensive computationally since the least-squares regressions are usually very quick to perform; again, see Beveridge and Joshi (2009).

In Section 5, we provide details of how the multiple regression enhancement is applied in our examples (see Section 5.3) and use it to demonstrate the usefulness of being able to obtain bounds. In particular, in some cases we see that the multiple regression enhancement leads to only very small changes in price, but can lead to substantial tightening in bounds.

4.3. **The Andersen–Broadie Method.** In the context of standard Bermudan derivatives, Andersen and Broadie (2004) developed a sub-Monte Carlo simulation upper-bound method that was both generic and accurate. By studying the optimal martingale, given by the Doob decomposition of the value process for the Bermudan, one can interpret it financially as the value of the following self-financing trading strategy: hold the underlying product where, if the exercise strategy says exercise, exercise the product and re-purchase the underlying product with one less exercise date, investing any residual cash in the numeraire asset; see, for example, Joshi (2007). The key idea behind the Andersen–Broadie method is that we can obtain an accurate approximation to the optimal martingale by holding the underlying product exercised according to an approximate exercise strategy, from the least-squares method for example.

We can apply the same idea for game options. In particular, we can interpret (3.2) as the value process for the following trading strategy. Fix an exercise strategy for B and take A's position in the underlying product exercised according to this strategy for B and the optimal strategy for A. At each exercise time, if the optimal strategy for A says exercise, then exercise the product and re-purchase with one less exercise time, otherwise continue holding the product. Note that this is the optimal hedge for B to hold against A in that however A exercises the product, whether default occurs, or whether B wants to exercise, B will always be covered; see, for example, Joshi (2007). Since we do not know the optimal strategy for A, we can apply the Andersen–Broadie method and approximate by using the least-squares strategy for A (or any other approximate strategy for that matter).

To make things concrete, we explain how to simulate each path when using Monte Carlo simulation to estimate (3.4) and (3.10). Clearly, the difficulty with the Andersen–Broadie method lies in simulating the approximation to the optimal martingales. We focus on simulating an approximation to $\tilde{M}(T_k, \tau_B^{\mathrm{LS}})$, since we can easily obtain an approximation for $\bar{M}(T_k, \tau_B^{\mathrm{LS}})$ from Section 3.1 by subtracting our least-squares price estimate. So for each path, we start by holding the underlying product from A's point of view exercised according to τ_A^{LS} (with τ_B^{LS} the strategy assumed for B), receiving the cash-flows generated by the product which we invest in the numeraire asset. At each exercise time, if neither strategy says exercise, the value of the approximation to the optimal martingale will equal the value of any cash-flows plus the value of the underlying product with one less exercise date. If either strategy says exercise or default occurs, the value of our approximation will equal the appropriate rebate plus the value of any intervening cash-flows.

In addition, if τ_A^{LS} says exercise, then we exercise and re-purchase the underlying product with one less exercise date. As such, we receive the corresponding exercise value, and have to pay for the new product. Any additional cash is handled by buying or selling the numeraire asset. If τ_B^{LS} says exercise or default occurs, we receive the corresponding rebate which we invest in the numeraire asset and cease hedging. If one is careful in keeping track of all the cash-flows, this is straightforward. The main subtlety arises in valuing the underlying product exercised according to τ_A^{LS} and τ_B^{LS}, which we need to value our approximation to the optimal martingale or to determine how much we have to pay to re-purchase the underlying product. As this problem is essentially the same as pricing at time zero, Monte Carlo simulation is required. That is, we have to run sub-Monte Carlo simulations to value the underlying product exercised according to our input exercise strategies.

The great advantage of this approach is that no further optimizations beyond those required for the least-squares method are necessary to develop the optimal martingale; once we have our two exercise strategies, we immediately have the associated approximation to the optimal martingale. The downside is that to value the hedge portfolio, we need to value the underlying product exercised according to the given strategies, and this generally requires sub-Monte Carlo simulations.

While the Andersen–Broadie method is intuitively appealing, the use of sub-Monte Carlo simulations means that our approximation to the optimal martingale is not strictly a martingale, and therefore can not be directly applied to the results obtained in Section 3. As such, before we can use this approach to calculate unbiased estimates of bounds for game options, we have to show that the use of sub-Monte Carlo simulations does not introduce bias in the wrong direction, which we now do. This corresponds to the fact that Andersen and Broadie had to show that their method led to an upwards bias in the single-exercisable case.

4.3.1. *Additive Bounds.* We show that the Monte Carlo errors of the sub-simulations only add to the upward bias in the additive approach. In particular, fix τ_A and τ_B, and let $\hat{M}(t, \tau_B)$ denote the value of the Andersen–Broadie approximate hedge. We can write

$$\hat{M}(T_j, \tau_B) = M(T_j, \tau_B) + \epsilon_{T_j},$$

where $M(T_j, \tau_B)$ is defined to be the exact value of the approximate hedge without approximation error at T_j, and ϵ_{T_j} is defined to be the accumulated Monte Carlo error along the path up to and including the time T_j. In addition, let $\underline{Z}_{T_j}$ denote the vector of random numbers used for the sub-simulations up to and including T_j, and $\mathcal{G}_{T_j} := \sigma(\underline{Z}_{T_j})$ denote the sigma-algebra generated by these.

Then we get the following.

Theorem 3. *If $\mathcal{F}_{T_m}$ and $\mathcal{G}_{T_m}$ are independent, then*

$$\mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - M(t_k, \tau_B)\right)\right] \leq \mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - \hat{M}(t_k, \tau_B)\right)\right]. \tag{4.1}$$

Proof. It is enough to prove that

$$\mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - M(t_k, \tau_B)\right) | \mathcal{F}_{T_m}\right] \leq \mathbb{E}\left[\max_{t_k \in \Xi^A} \left(f_0(t_k, \tau_B) - \hat{M}(t_k, \tau_B)\right) | \mathcal{F}_{T_m}\right],$$

since (4.1) then follows by the Tower Law. As $\hat{M}(T_j, \tau_B) = M(T_j, \tau_B) + \epsilon_{T_j}$ and given $\mathcal{F}_{T_m}$ everything is known except the ϵ's, this is equivalent to showing

$$\mathbb{E}\left[\max_i (\alpha_i + \gamma_i)\right] \geq \max_i (\alpha_i), \tag{4.2}$$

for a sequence of real numbers $\{\alpha_i\}$, and a sequence of random variables $\{\gamma_i\}$ with $\mathbb{E}[\gamma_i] = 0$. The γ_i's represent the accumulated Monte Carlo errors in $\hat{M}(T_j, \tau_B)$, ϵ_{T_j}.

Now, let k be such that $\alpha_k = \max_i(\alpha_i)$. Then

$$\begin{aligned} \max_i (\alpha_i) &= \alpha_k, \\ &= \mathbb{E}[\alpha_k + \gamma_k], \end{aligned}$$

since $\mathbb{E}[\gamma_k] = 0$. Now since we clearly have

$$\mathbb{E}[\alpha_k + \gamma_k] \leq \mathbb{E}\left[\max_i (\alpha_i + \gamma_i)\right],$$

we have proved (4.2), which proves our result.

$\square$

4.3.2. *Multiplicative Bounds.* In the following, we stick with the approximate martingale described in Section 4.3, rather than use the alternative approach suggested by Joshi (2007) for using the Andersen–Broadie method with the multiplicative approach. In order to do this we make the rather strong assumption that the approximate hedge is strictly positive. Although this seems reasonable, the cumulative re-hedging error could potentially be an issue. In particular, each time our approximate strategy says exercise when the estimated continuation value is above the exercise value, we have to sell the numeraire asset to fund our re-purchase of the underlying product. Further, there is generally nothing to ensure that the cumulative re-hedging error does not force the value of the hedge portfolio below zero.

Alternatively, we can use the re-investment strategy for the optimal hedge advocated by Joshi (2007). Under this strategy, all cash flows received are reinvested in the underlying product rather than in the numeraire asset. Therefore, since we will always be investing positive cash-flows into a product with positive value by assumption, our approximate hedge will remain positive and it can easily be shown using similar arguments to those of Joshi (2007) that the Andersen–Broadie method does not introduce bias in the wrong direction. This holds regardless of the size of the re-hedging errors; all that is required is that we must obtain strictly positive values for the underlying product, which can be ensured by using a suitable number of sub-simulation paths.

In all practical examples considered, the cumulative re-hedging error build up described above for the additive reinvestment strategy was not a problem, indicating that the method is indeed of practical value. Using the same martingale for the two methods has the virtue that that they can be run simultaneously with only a very small amount of additional work, since along each path they require very similar numbers. We also feel that is useful for the reader to see a proof that the additive reinvestment strategy is applicable for the multiplicative bounds since this is new even in the single exercisable case.

Again fix τ_A and τ_B, and consider the notation from Section 4.3.1 with some minor alterations. In particular, change $\hat{M}(t, \tau_B)$ to denote the value of the Andersen–Broadie approximate hedge starting with the initial product. In addition, let $\mathcal{H}_{T_j}$ denote the sigma-algebra generated by the random numbers of the T_j sub-simulation. Then we get the following.

Theorem 4. *If $\mathcal{F}_{T_m}$ and $\mathcal{G}_{T_m}$ are independent, $\mathcal{H}_{T_j}$ and $\mathcal{H}_{T_k}$ are independent for all $j \neq k$, and $\hat{M}(t, \tau_B)$ is strictly positive, then*

$$\mathbb{E}\left[\max_{T_k \in \Xi^A}\left(\frac{f_0(T_k, \tau_B)}{M(T_k, \tau_B)}\right) M(T_m, \tau_B)\right] \leq \mathbb{E}\left[\max_{T_k \in \Xi^A}\left(\frac{f_0(T_k, \tau_B)}{\hat{M}(T_k, \tau_B)}\right) \hat{M}(T_m, \tau_B)\right].$$

Proof. By the Tower Law, it is enough to prove that

$$\mathbb{E}\left[\max_{T_k \in \Xi^A}\left(\frac{f_0(T_k, \tau_B)}{M(T_k, \tau_B)}\right) M(T_m, \tau_B) | \mathcal{F}_{T_m}\right] \leq \mathbb{E}\left[\max_{T_k \in \Xi^A}\left(\frac{f_0(T_k, \tau_B)}{\hat{M}(T_k, \tau_B)}\right) \hat{M}(T_m, \tau_B) | \mathcal{F}_{T_m}\right].$$

Since $\hat{M}(T_j, \tau_B) = M(T_j, \tau_B) + \epsilon_{T_j}$ and given $\mathcal{F}_{T_m}$ everything is known except the ϵ's, this is equivalent to showing

$$\mathbb{E}\left[\max_{i\in\{1,\ldots,m\}}\left(\frac{\alpha_i}{\beta_i+\gamma_i}\right)(\beta_m+\gamma_m)\right] \geq \max_{i\in\{1,\ldots,m\}}\left(\frac{\alpha_i}{\beta_i}\right)(\beta_m),$$

where $\{\alpha_i\}$ and $\{\beta_i\}$ are sequences of positive real numbers, and

$$\gamma_i := \pi_0 + \ldots + \pi_i,$$

with the π_i being independent random variables satisfying $\mathbb{E}[\pi_i] = 0$. The π_i's represent the Monte Carlo errors for the individual sub-Monte Carlo simulations, and are independent by the second independence assumption. The α_i's and β_i's represent the game option payoff and martingale value (without Monte Carlo error) respectively, and are therefore positive by assumption.

Now,

$$\begin{aligned}
\text{LHS} &= \mathbb{E}\left[\max_{i\in\{1,\ldots,m\}}\left(\frac{\alpha_i}{\beta_i+\gamma_i}(\beta_m+\gamma_m)\right)\right], \\
&\geq \max_{i\in\{1,\ldots,m\}}\left(\mathbb{E}\left[\frac{\alpha_i}{\beta_i+\gamma_i}(\beta_m+\gamma_m)\right]\right).
\end{aligned}$$

Since we can write,

$$\frac{(\beta_m+\gamma_m)}{\beta_i+\gamma_i} = 1 + \frac{\beta_m+\gamma_m-(\beta_i+\gamma_i)}{\beta_i+\gamma_i},$$

then

$$\begin{aligned}
\text{LHS} &\geq \max_{i\in\{1,\ldots,m\}}\left(\alpha_i\mathbb{E}\left[1+\frac{\beta_m+\gamma_m-(\beta_i+\gamma_i)}{\beta_i+\gamma_i}\right]\right), \\
&= \max_{i\in\{1,\ldots,m\}}\left(\alpha_i\left(1+\mathbb{E}\left[\beta_m+\gamma_m-(\beta_i+\gamma_i)\right]\mathbb{E}\left[\frac{1}{\beta_i+\gamma_i}\right]\right)\right),
\end{aligned}$$

where the final line follows by the independence of the π_i's. Simplifying further using Jensen's inequality together with the assumption that $\hat{M}$ is strictly positive so $(\beta_i+\gamma_i) > 0$,

$$\begin{aligned}
\text{LHS} &\geq \max_{i\in\{1,\ldots,m\}}\left(\alpha_i\left(1+(\beta_m-\beta_i)\frac{1}{\mathbb{E}\left[\beta_i+\gamma_i\right]}\right)\right), \\
&= \max_{i\in\{1,\ldots,m\}}\left(\frac{\alpha_i\beta_m}{\beta_i}\right), \\
&= \max_{i\in\{1,\ldots,m\}}\left(\frac{\alpha_i}{\beta_i}\right)\beta_m,
\end{aligned}$$

and we are done. □

5. Results

We consider the pricing of convertible bonds, showing the usefulness of our new results. Examples are built from those of Lvov et al. (2004a), Lvov (2005), and Crépey and Rahal (2009), representing realistic and challenging cases. We study two models; the first is a typical model used by Lvov et al. (2004a); the second is significantly more complicated since it allows stochastic volatility and variable default intensity.

5.1. **Set-up.** We now describe the two models used in our tests. For each model, we use the risk-neutral measure for pricing, and specify all dynamics in this measure.

5.1.1. *Model 1.* We use a two-factor model with a stochastic stock price and stochastic interest rates. Prior to default, the stock price is assumed to follow geometric Brownian motion. That is,

$$\frac{d\tilde{S}(t)}{\tilde{S}(t)} = \mu(t)dt + \sigma dW_1(t).$$

In addition, the short rate, denoted $r(t)$, is assumed to follow a CIR process (see Cox et al. (1985)),

$$dr(t) = \vartheta(\theta - r(t))dt + \upsilon\sqrt{r(t)}dW_2(t),$$

where $r(0) > 0$, we assume that $2\vartheta\theta > \upsilon^2$ to ensure positivity, and the correlation between $W_1(t)$ and $W_2(t)$ is given by $\rho_{1,2}$. Also, we assume that the default intensity is constant and given by λ, and that upon default at time t, the value of the stock price is given by $\delta\tilde{S}(t)$. With this notation, we have that

$$\mu(t) = r(t) - q + (1-\delta)\lambda,$$

where q denotes the dividend yield.

Since the CIR process used for the short rate satisfies the positivity constraint, we use the discretisation scheme of Kahl and Jäckel (2006), to simulate our state variables. In particular, we avoid the use of more complicated schemes that are primarily needed when the positivity constraint fails.

5.1.2. *Model 2.* We use a Heston model for the stock together with non-constant default intensity and stochastic interest rates. In particular, prior to default the stock price is assumed to follow

$$\frac{d\tilde{S}(t)}{\tilde{S}(t)} = \mu(\tilde{S}(t), t)dt + \sqrt{\nu(t)}dW_1(t),$$

where

$$d\nu(t) = \kappa(\alpha - \nu(t))dt + \epsilon\sqrt{\nu(t)}dW_3(t),$$

so the instantaneous variance is stochastic and follows a CIR process. We again assume that the short rate, $r(t)$, follows a CIR process,

$$dr(t) = \vartheta(\theta - r(t))dt + \upsilon\sqrt{r(t)}dW_2(t),$$

where $r(0) > 0$ and we assume the positivity constraint $2\vartheta\theta > \upsilon^2$. In addition, we denote the correlation between $W_i(t)$ and $W_j(t)$ by $\rho_{i,j}$.

We assume that the default intensity, $\lambda(\tilde{S}(t), t)$, is given by

$$0.03\left(\frac{\tilde{S}(0)}{\tilde{S}(t)}\right)^{1.2},$$

and that upon default at time t, the value of the stock price is given by $\delta\tilde{S}(t)$. Again,

$$\mu(\tilde{S}(t),t) = r(t) - q + (1-\delta)\lambda(\tilde{S}(t),t),$$

where q denotes the dividend yield.

To evolve the state variables of this model we use the popular QE algorithm introduced by Andersen (2008).

5.2. **Convertible Bond Pay-Off.** We take the continuously compounding money market account as numeraire, and as such take $N(t) = \exp\left(\int_0^t r(s)ds\right)$. In addition, assume that we have a fixed set of dates, $\mathcal{T} = \{t_1, \dots, t_n\}$, where all cash-flows and exercise times occur, with

$$0 < t_1 < \dots < t_n = T.$$

We use t_i rather than T_j (and correspondingly n rather than m) to emphasize that cash-flow times may not be exercise times. For convertible bonds, the pay-off function f is given by

$$f_t(\tau_i,\tau_h) = \sum_{k=\eta(t)}^{\eta(\min(\tau_i,\tau_h,\tau_d))-1} D_k + \mathrm{I}_{\tau_h\wedge\tau_i<\tau_d}\left(\mathrm{I}_{\tau_h\le\tau_i}H_{\tau_h} + \mathrm{I}_{\tau_i<\tau_h}C_{\tau_i}\right) + \mathrm{I}_{\tau_d\le\tau_i\wedge\tau_h}R_{t_{\eta(\tau_d)}},$$

where $\eta(t)$ is a right-continuous function giving the index of the next date in $\mathcal{T}$ as of time t, and we have replaced A and B with h and i, representing the holder and issuer of the convertible bond. As such, D_k gives the deflated coupon payments occurring at each time t_k, H_{t_k} gives the deflated exercise value if the holder exercises at t_k and C_{t_k} that of the issuer. Note that the holder has priority over the issuer when exercising the bond. In addition, R_{t_k} gives the rebate received by the holder if default occurs between t_{k-1} and t_k and we assume that the default payment occurs at the next time in $\mathcal{T}$ from where default occurred. Further, for our examples we assume the following forms for the exercise values and exercise dates,

- $\tilde{D}_j = D_j N(t_j)$ denotes the non-deflated coupon payments. We assume 16 exercise times per year, and half-yearly coupon payments, so $\tilde{D}_j$ will be zero except when $j = 8k$ for $k \in \mathbb{N}$, at which times it will equal the size of the coupon, which we assume to be constant.
- $H_{t_n} = C_{t_n} = \max(a\tilde{S}(t_n), F + \tilde{D}_n)/N(t_n)$, where a denotes the conversion ratio and F the face value for the convertible bond.
- $H_{t_j} = \left(a\tilde{S}(t_j)\mathrm{I}_{t_j\neq t_p} + \max(a\tilde{S}(t_j), P + AI_j)\mathrm{I}_{t_j=t_p}\right)/N(t_j)$, where AI_j denotes the accrued interest at t_j. We assume a single time at which the bond can be put by the holder, which we denote by t_p, and allow the bond to be converted into stock by the holder at any time in $\mathcal{T}$. As such, the exercise value of the holder is given by the maximum of the conversion value and the put value, when it is possible to put.
- $C_{t_j} = (C + AI_j)/N(t_j)$, so the issuer's exercise value is simply given by the call value.
- $R_{t_{\eta(\tau_d)}} = \max(a\delta\tilde{S}(\tau_d), \delta^* F)/N(t_{\eta(\tau_d)})$, where δ^* gives the proportion of the face value that can be recovered upon default. As such, the rebate received upon default is the maximum of the conversion value and the reduced face value.

Note that we have defined the exercise value at maturity in such a way that we can assume that the contract is always exercised then if not before which ensures finiteness of exercise times.

In terms of exercise dates, since the bond can be converted into stock by the holder at any time in $\mathcal{T}$,

$$\Xi^h = \{t_1, \ldots, t_n\}.$$

In contrast, for the issuer we assume a more complicated structure. In particular, with ϑ a random variable, we let

$$\Xi^i = \{t_\vartheta, t_{\vartheta+1}, t_{\vartheta+2}, \ldots, t_n\}.$$

We consider two choices for ϑ, corresponding to two different forms of call protection. The first is a fixed lockout period as used by Lvov et al. (2004a), so

$$\vartheta_1 = c.$$

In addition, we also consider a highly path-dependent form of call protection used by Crépey and Rahal (2009), and let

$$\vartheta_2 = \min\left\{ j : \sum_{k=1}^{d} U_k \geq l \right\} \wedge t_n.$$

Here U is a vector of indicator functions of the events $\tilde{S}(t_j) \geq \bar{S}$ for the last d dates in $\mathcal{T}$ from the current time. So here ϑ_2 represents the first time such that the stock price is above $\bar{S}$ on l of the last d checking dates. As explained by Crépey and Rahal (2009), this form of call protection is commonly used in practice, but it is strongly path-dependent and rules out the use of lattice methods for large values of d.

Unless specified otherwise, the pricing algorithm described by Lvov et al. (2004a) is used to calculate a least-squares price, for which we obtain bounds using our new approach. When calculating least-squares prices, to both develop our approximate exercise strategy and calculate an out-of-sample price, we follow Lvov et al. (2004a) and use 100,000 paths for 3 and 5 year maturities, 150,000 paths for 7 and 10 year maturities, and 200,000 paths when the maturity is 15 years. However, due to memory constraints we use 150,000 paths for the 15 year maturity with Model 2 and ϑ_2. For calculating bounds, 1000 paths are used for both the outer and inner simulations. We use *Mersenne Twister* pseudo-random numbers throughout, and evolve our state variables between each exercise date in one step.

Rather than mimic the comprehensive tests from Lvov et al. (2004a) and Lvov (2005), we focus on the contract terms and model parameters from these tests where the least-squares method struggled the most, as well as imposing additional highly path-dependent call protection used by Crépey and Rahal (2009) and an enhanced model.

In particular, we use the model parameters in Tables 5.1 and 5.2 throughout, and consider each set of contract terms in Table 5.3 supplemented with the call protection terms in Table 5.4. We use Model 1 with ϑ_1 to apply our new methods to the toughest examples from Lvov (2005), and consider Model 2 with ϑ_2 to provide a more realistic and challenging example. While we refer the reader to Lvov (2005) for further discussion, we do stress that the model parameters and contract terms were chosen to closely represent those often seen in practice. In particular, after pricing over 10,000 different convertible bonds, it was the parameters for Model 1 and contract terms with ϑ_1 which we consider that produced the greatest errors in the least-squares prices.

For basis functions, when using Model 1 with ϑ_1 we use the first three powers of the stock price and the interest rate, their product, and a constant; these are the same basis functions as used by Lvov et al. (2004a). In contrast, due to the increased complexity when using Model 2 with ϑ_2, we use quadratic polynomials in the stock price, short rate, instantaneous variance of the stock (denoted by $\nu(t)$), and

$$\sum_{k=1}^{d} U_k.$$

The final variable is used to encapsulate the path-dependent nature of the call protection. We have chosen this variable since it interacts well with the multiple regression enchancement which we will discuss below, and we shall see that it leads to tight bounds. We note that an alternate more complex approach to the path-dependence is discussed in Crépey and Rahal (2009), however, we shall see that this variable in conjunction with multiple regression is sufficient.

5.3. **Multiple Regression Enhancement.** We describe how to apply the multiple regression enhancement for the examples considered. For call dates, we partition the stock price into four regions:

- $\Pi_1^j = (C + AI_j, \infty)$,
- $\Pi_2^j = (C + AI_j - 20, C + AI_j]$,
- $\Pi_3^j = (P + AI_j, C + AI_j - 20]$,
- $\Pi_4^j = (-\infty, P + AI_j]$,

and estimate the continuation value conditional on $\tilde{S}(t_j) \in \Pi_k^j$ for each k. Since we assume the dividend yield is zero, the continuation value cannot be smaller than the conversion value; see Ingersoll (1977). As such, we know the issuer will call if $\tilde{S}(t_j) \in \Pi_1^j$ and therefore we do not need to estimate the continuation value to determine the exercise strategy in that region. For the remaining three regions, we use separate least-squares regressions. To estimate the conditional expectation for region k, we perform a least-squares regression as described by Lvov et al. (2004a), but only include points satisfying $\tilde{S}(t_j) \in \Pi_k^j$ where it is possible exercise may occur by at least one party. To determine our approximate exercise strategy, we first determine which region $\tilde{S}(t_j)$ is in, and then use our estimated continuation value conditional on $\tilde{S}(t_j)$ being in that region.

Note that at non-call and non-put dates, that is where only conversion is possible, we know that the bond should not be converted, and we do not need to perform any regressions to estimate continuation values. This follows since the continuation value must be at least as big as the conversion value; again, see Ingersoll (1977).

5.4. **Calculating Additive and Multiplicative Bounds.** When using the Andersen–Broadie method, there is a natural decomposition that can be applied to additive bounds; we can naturally write an additive bound as the sum of the least-squares price and the distance between this price and the bound. In particular, fix τ_A and τ_B, and consider the approximate hedge from Section 4.3.1

together with (3.1),

$$\begin{aligned} V(0) &\leq \mathbb{E}\left[\max_{t_k\in\Xi^A}\Big(f_0(t_k,\tau_B)-\hat{M}(t_k,\tau_B)\Big)\right], \qquad (5.1)\\ &= \mathbb{E}\left[\max_{t_k\in\Xi^A}\Big(f_0(t_k,\tau_B)-(\hat{M}(t_k,\tau_B)+\mathbb{E}\left[f_0(\tau_A,\tau_B)\right]+\epsilon_0)\Big)+\mathbb{E}\left[f_0(\tau_A,\tau_B)\right]+\epsilon_0\right],\\ &= \mathbb{E}\left[\max_{t_k\in\Xi^A}\Big(f_0(t_k,\tau_B)-(\hat{M}(t_k,\tau_B)+\mathbb{E}\left[f_0(\tau_A,\tau_B)\right]+\epsilon_0)\Big)\right]+\mathbb{E}\left[f_0(\tau_A,\tau_B)\right]. \qquad (5.2)\end{aligned}$$

Rather than calculate (5.1) directly, we can use the modified hedge which starts with the initial product and estimate the two terms in (5.2) separately. Calculating the first term with Monte Carlo simulation gives an estimate of the distance between the exact least-squares price and the corresponding bound, while the second term is just the least-squares price, which can be calculated using the least-squares method. When reporting bounds for the additive method, we provide estimates of the distance between the bound and the true least-squares price corresponding to the first term in (5.2). This has the advantage that the Monte Carlo error does not dominate the difference between the least-squares prices and the corresponding bounds, allowing for more meaningful analysis.

Unfortunately, no such decomposition can naturally be applied to the multiplicative method, and we report bounds as prices, not distances from the least-squares price.

5.5. **Numerical Results.** Numerical results without the multiple regression enhancement are given in Tables 5.5 and 5.7 for Model 1 with ϑ_1 and Model 2 with ϑ_2 respectively, while those with the multiple regression enhancement are in Tables 5.6 and 5.8. These results illustrate a number of key points. First, for challenging and realistic examples, our new (additive) approach allows the calculation of tight bounds for convertible bonds. In particular, the greatest distance between the lower and upper bounds was 0.51 without the multiple regression enhancement, and 0.17 with it; these bounds are less than 0.4% of the bond's value in each case, with much tighter bounds obtained when using the multiple regression enhancement.

Second, the additive approach clearly outperforms the multiplicative approach in terms of precision. The standard errors of the multiplicative bounds are always at least ten times that of the additive bounds, where the standard error of the additive bounds are given by the square-root of the sum of the squares of the standard errors of the least-squares price and the corresponding distance. This occurs for two reasons. The main reason is the additive approach can use the least-squares price, for which we can use a large number of paths and therefore obtain low standard errors relatively efficiently, as a starting point and calculate the bounds relative to this price; see (5.2). In contrast, this is not possible for the multiplicative approach, and we must essentially price the bond as we do our outer simulation, severely limiting the number of paths we can use and therefore the precision we can obtain. To a lesser extent, from Remark 1, the additive approach gives zero variance when the optimal martingale is used, but there is no corresponding result for the multiplicative approach (also see Chen and Glasserman (2007)).

Third, the multiple regression enhancement can lead to significant reductions in the bounds. For instance, in Example 4 with Model 1 and ϑ_1 the multiple regression enhancement reduces the distance between the lower and upper bounds by a factor of 3.5. Similarly, in Example 3 with

$S(0)$	σ	q	$r(0)$	ϑ	θ	υ	$\rho_{1,2}$	λ	δ
80	0.6	0	0.1	1.0	0.05	0.2	0	0.03	0.3

TABLE 5.1. Parameters for Model 1.

$S(0)$	$\nu(0)$	κ	α	ϵ	q	$r(0)$	ϑ	θ	υ	$\rho_{1,2}$	$\rho_{1,3}$	$\rho_{2,3}$	δ
80	0.1	1.0	0.1	0.5	0	0.1	1.0	0.05	0.2	0	-0.7	0	0.3

TABLE 5.2. Parameters for Model 2.

Example	T	t_p	C	P	a	$\tilde{D}_j$	F	δ^*
1	3	2	150	100	1.2	$4\mathrm{I}_{\{j=8k,k\in\mathbb{N}\}}$	100	0.3
2	5	2	150	110	1.2	$4\mathrm{I}_{\{j=8k,k\in\mathbb{N}\}}$	100	0.3
3	7	2	160	110	1.2	$4\mathrm{I}_{\{j=8k,k\in\mathbb{N}\}}$	100	0.3
4	10	2	160	110	1.2	$4\mathrm{I}_{\{j=8k,k\in\mathbb{N}\}}$	100	0.3
5	15	4	160	110	1.2	$4\mathrm{I}_{\{j=8k,k\in\mathbb{N}\}}$	100	0.3

TABLE 5.3. Contract terms considered.

c	d	l	$\bar{S}$
16	16	12	83

TABLE 5.4. Call protection terms for ϑ_1 and ϑ_2.

Model 2 and ϑ_2 the distance between upper and lower bounds is reduced by a factor of 5.1. It is interesting to note that in these cases the least-squares prices are very similar. In particular, in the first instance the least-squares prices only differ by 0.15, about one standard error, illustrating the importance of being able to calculate bounds to assess the accuracy of the least-squares prices.

Fourth, the distance between upper and lower bounds is greater in the more difficult case of Model 2 with the highly path-dependent form of call protection, as is to be expected. However, as mentioned above, the bounds obtained are still tight, especially when the multiple regression enhancement is used.

Finally, the multiplicative lower bound is above the corresponding upper bound with and without the multiple regression enhancement for Example 3 with Model 1 and ϑ_1. However, we believe that this is due to Monte Carlo error, and an increase in the number of outer paths would see this discrepancy disappear.

6. Conclusion

We have introduced a new method which generalises the duality approaches of Rogers (2002), Haugh and Kogan (2004) and Jamshidian (2004) to the case of game options. In addition, we have shown that the Andersen–Broadie method can be applied as a generic way to implement the new results. Combining these new findings, we have a generic and accurate approach to calculating bounds for game options using Monte Carlo simulation. The new results are applied to convertible bonds, and we demonstrate that it is indeed possible to obtain tight bounds in challenging and realistic examples.

		Additive Bounds		Multiplicative Bounds	
Example	LS	Lower	Upper	Lower	Upper
1	132.57 (0.17)	-0.065 (0.012)	0.022 (0.005)	131.64 (1.79)	132.36 (1.77)
2	137.19 (0.17)	-0.120 (0.017)	0.019 (0.005)	137.23 (1.80)	137.36 (1.82)
3	139.00 (0.14)	-0.166 (0.022)	0.031 (0.008)	137.89 (1.61)	137.59 (1.65)
4	139.73 (0.14)	-0.242 (0.028)	0.074 (0.014)	138.52 (1.75)	140.86 (1.76)
5	141.19 (0.12)	-0.175 (0.023)	0.044 (0.011)	139.86 (1.80)	141.94 (1.80)

TABLE 5.5. Bounds for convertible bounds using Model 1 and ϑ_1 without the multiple regression enhancement. Non-bracketed numbers give prices and bracketed numbers give the corresponding standard errors. The LS column gives the least-squares price.

		Additive Bounds		Multiplicative Bounds	
Example	LS	Lower	Upper	Lower	Upper
1	132.14 (0.17)	-0.017 (0.004)	0.004 (0.002)	133.30 (1.77)	133.43 (1.77)
2	137.09 (0.17)	-0.014 (0.003)	0.009 (0.003)	137.13 (1.80)	137.48 (1.82)
3	138.86 (0.14)	-0.027 (0.005)	0.028 (0.008)	137.93 (1.60)	137.35 (1.65)
4	139.58 (0.14)	-0.026 (0.006)	0.063 (0.013)	138.88 (1.75)	140.45 (1.77)
5	141.04 (0.12)	-0.026 (0.006)	0.035 (0.010)	139.89 (1.79)	141.69 (1.80)

TABLE 5.6. Bounds for convertible bounds using Model 1 and ϑ_1 with the multiple regression enhancement. Non-bracketed numbers give prices and bracketed numbers give the corresponding standard errors. The LS column gives the least-squares price.

		Additive Bounds		Multiplicative Bounds	
Example	LS	Lower	Upper	Lower	Upper
1	119.87 (0.09)	-0.367 (0.028)	0.088 (0.017)	119.32 (1.00)	119.81 (1.01)
2	126.31 (0.10)	-0.396 (0.022)	0.049 (0.012)	126.47 (1.09)	127.14 (1.08)
3	129.24 (0.10)	-0.479 (0.027)	0.032 (0.009)	130.14 (1.18)	131.09 (1.20)
4	129.04 (0.10)	-0.353 (0.022)	0.070 (0.015)	128.51 (1.25)	129.79 (1.26)
5	129.52 (0.11)	-0.416 (0.024)	0.029 (0.008)	129.15 (1.47)	129.47 (1.52)

TABLE 5.7. Bounds for convertible bounds using Model 2 and ϑ_2 without the multiple regression enhancement. Non-bracketed numbers give prices and bracketed numbers give the corresponding standard errors. The LS column gives the least-squares price.

		Additive Bounds		Multiplicative Bounds	
Example	LS	Lower	Upper	Lower	Upper
1	119.57 (0.09)	-0.049 (0.010)	0.056 (0.015)	119.36 (0.99)	119.62 (1.01)
2	126.09 (0.10)	-0.136 (0.012)	0.004 (0.002)	126.63 (1.07)	126.93 (1.09)
3	128.85 (0.10)	-0.093 (0.008)	0.007 (0.004)	130.17 (1.18)	130.60 (1.21)
4	128.85 (0.10)	-0.120 (0.010)	0.064 (0.014)	128.53 (1.25)	129.61 (1.27)
5	129.30 (0.11)	-0.159 (0.014)	0.029 (0.008)	129.20 (1.47)	129.21 (1.53)

TABLE 5.8. Bounds for convertible bounds using Model 2 and ϑ_2 with the multiple regression enhancement. Non-bracketed numbers give prices and bracketed numbers give the corresponding standard errors. The LS column gives the least-squares price.

References

Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *Journal of Computational Finance*, *11*, 1-42.

Andersen, L., & Broadie, M. (2004). A primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, *50*, 1222-1234.

Andersen, L., & Buffum, D. (2002). *Calibration and implementation of convertible bond models.* (http://ssrn.com/abstract=355308)

Beveridge, C. J., & Joshi, M. S. (2009). *Practical policy iteration: generic methods for obtaining rapid and tight bounds for Bermudan exotic derivatives using Monte Carlo simulation.* (http://ssrn.com/abstract=1331904)

Bielecki, T. R., Crépey, S., Jeanblanc, M., & Rutkowski, M. (2008). Arbitrage pricing of defaultable game options with applications to convertible bonds. *Quantitative Finance*, *8*, 795-810.

Carrière, J. F. (1996). Valuation of the early-exercise price for options using simulation and nonparametric regression. *Insurance: Mathematics and Economics*, *19*, 19-30.

Chassagneux, J.-F., & Crépey, S. (2009). *Doubly reflected BSDEs with call protection and their approximation.* (http://grozny.maths.univ-evry.fr/pages_perso/crepey/papers/III-3.pdf)

Chen, N., & Glasserman, P. (2007). Additive and multiplicative duals for American option pricing. *Finance and Stochastics*, *11*, 153-179.

Cox, J., Ingersoll, J. E., & Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, *53*, 385-407.

Crépey, S., & Rahal, A. (2009). *Pricing convertible bonds with call protection.* (http://grozny.maths.univ-evry.fr/pages_perso/crepey/papers/II-A-2.pdf)

Davis, M. H. A., & Karatzas, I. (1994). A deterministic approach to optimal stopping, with applications. In F. P. Kelly (Ed.), *Probability, statistics and optimisation: a tribute to Peter Whittle* (p. 455-466). New York: John Wiley and Sons Ltd.

Fries, C. P. (2005). *Foresight bias and suboptimality correction in Monte-Carlo pricing of options with early exercise: classification, calculation and removal.* (http://ssrn.com/abstract=839105)

Geman, H., El Karoui, N., & Rochet, J. C. (1995). Changes of numeraire, changes of probability measures and pricing of options. *Journal of Applied Probability*, *32*, 443-458.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering.* New York: Springer.

Grimwood, R., & Hodges, S. (2002). *The valuation of convertible bonds: a study of alternative pricing models.* (http://www2.warwick.ac.uk/fac/soc/wbs/research/wfri/wpaperseries/pp_02.121.pdf)

Haugh, M., & Kogan, L. (2004). Pricing american options: A duality approach. *Operations Research*, *52*, 258-270.

Ingersoll, J. (1977). A contingent-claims valuation of convertible securities. *Journal of Financial Economics*, *4*, 289-322.

Jamshidian, F. (2004). *Numeraire-invariant option pricing and american, bermudan, and trigger stream rollover.* (http://staff.science.uva.nl/˜spreij/stieltjes/jamshidian.pdf)

Joshi, M. S. (2006). *Monte Carlo bounds for callable products with non-analytic break costs.* (http://ssrn.com/abstract=907407)

Joshi, M. S. (2007). A simple derivation of and improvements to Jamshidian's and Rogers' upper bound methods for Bermudan options. *Applied Mathematical Finance*, *14*, 197-205.

Kahl, C., & Jäckel, P. (2006). Fast strong approximation Monte-Carlo schemes for stochastic volatility models. *Quantitative Finance*, *6*, 513-536.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: a simple least squares approach. *The Review of Financial Studies*, *14*, 113-147.

Lvov, D. (2005). *Monte Carlo methods for pricing and hedging: applications to Bermudan swaptions and convertible bonds.* Unpublished doctoral dissertation, University of Reading.

Lvov, D., Yigitbasioglu, A. B., & El Bachir, N. (2004a). Pricing convertible bonds by simulation. In *Second IASTED international conference* (p. 259-264). Cambridge, MA, USA.

Lvov, D., Yigitbasioglu, A. B., & El Bachir, N. (2004b). *Pricing convertible bonds by simulation.*

Rogers, L. C. G. (2002). Monte Carlo valuation of American options. *Mathematical Finance*, *12*, 271-286.

Takahashi, A., Kobayashi, T., & Nakagawa, N. (2001). Pricing convertible bonds with default risk. *Journal of Fixed Income*, *11*, 20-29.

Tsiveriotis, K., & Fernandes, C. (1998). Valuing convertible bonds with credit risk. *Journal of Fixed Income*, *8*, 95-102.

Centre for Actuarial Studies, Dept of Economics, University of Melbourne, Victoria 3010, Australia