

Delta Least Squares Monte Carlo Pricing of American Options

Peter Pommergård Lind

Rolf Poulsen

Aalborg University

University of Copenhagen

ppl@business.aau.dk

rolf@math.ku.dk

mrppl.github.io

April 9, 2024

Abstract

We present a new simulation-based American option pricing method, *Delta least squares Monte Carlo* (Delta LSM). Whereas the classical LSM method from Longstaff & Schwartz (2001) uses only the discounted payoff to learn the continuation value, Delta LSM uses both the discounted payoff and its derivative (Delta) to estimate regression coefficients. The Delta LSM is straightforward to implement and comes at little extra numerical cost. It is quite literally an add-on to the LSM method. Our numerical experiments show that irrespective of your speed/safety preference – and robustly across market scenarios – Delta LSM gives a marked improvement over classical LSM.

Keywords: American option pricing, Longstaff-Schwartz least squares Monte Carlo, Delta regularization.

1 Introduction

Most exchange traded options are American-style, i.e. they give their holders early exercise rights. However, even for the simplest product (an American put option) under minimal realistic market conditions (interest rate $>$ dividend yield ≥ 0) and the most sanitized mathematical model (Geometric Brownian motion as dynamics for the underlying and frictionless continuous trading), there is no closed-form expression for neither the "right" (=unique arbitrage-free) price, nor for the "right" (=optimal) exercise strategy. A plethora of articles have developed specific-to-possibly-general theoretical representations and numerical methods, many of which are very clever and efficient; to our knowledge Andersen, Lake & Offengenden (2016) is the gold standard. That said, most production implementations in the financial industry use a general-to-specific approach with an outset in the least squares Monte Carlo method (LSM) presented in Longstaff & Schwartz (2001), because of its, well, generality. Again, a full and fair review of the underlying body of literature is beyond the scope of this paper – and arguably its authors. So we will simply refer to Becker, Cheridito, Jentzen & Welte (2021) and Lind (2022) for treatments of the state-of-the-art using neural network regression to handle multi-dimensionality but at the cost of analytic tractability such as nonexistent theoretical convergence results.

Our aim with this paper is to introduce, describe, and investigate an analytical tractable method that can be implemented as an add-on to any LSM-type algorithm. Specifically, the method (section 2) relies on the Delta regularization proposed in Huge & Savine (2020) and applied to European option problems in Frandsen, Pedersen & Poulsen (2022). The Delta LSM extends the classical least squares Monte Carlo method¹ from Longstaff & Schwartz (2001) by also using the derivative of the discounted cash flow to find the optimal stopping strategy. The additional algorithmic complication and computational cost of including the derivative of the discounted payoff in the loss function amount only to that of adding two matrices. This is a relatively small cost since the heavy computation in

¹We shall write 'LSM' as shorthand, but 'LSM method' in readable text to improve the flow. By interpreting M as coming from 'Monte' rather than 'method', this technically isn't a pleonasm. The first letter equivalence of 'least squares' and 'Longstaff & Schwartz' is a mnemonic serendipity.

least-squares lies in the matrix inversion. Our experimental results (section 3) show that in return for this small additional complexity, the Delta LSM method delivers – robustly across market scenarios – better results than LSM irrespective of whether your preference is for speed (the in-sample bias and lower bound valuation is improved, particularly for small samples) or for safety (it gives much tighter upper bounds when run through the Primal-Dual-algorithm from Andersen & Broadie (2004)).

2 Delta LSM for American options

The least squares Monte Carlo method as presented in Longstaff & Schwartz (2001) is ubiquitous for the valuation of anything involving stopping decisions, be that American-style financial contracts or real options; *the American option* in our language. Structurally, the method can be split up into four parts:

1. Simulation of paths.
2. Backward recursion to estimate regression coefficients (i.e. parameters minimizing some loss function – here of least squares type as suggested by the name) determining the optimal exercise decision at each time point.
3. Forward pass: Simulate paths as in (but independently of) step 1 and use the regression coefficients from step 2 to decide when to exercise the American option – hopefully close to optimally.
4. Valuation: Average over the discounted payoffs from the decisions in step 3 to estimate the American option price.

The Delta LSM method, which we shall describe in more detail in the following subsections, is an extension to step 2 where the loss function is regularized using the Delta of the continuation value of the American option, i.e. the continuation value's sensitivity wrt. the underlying, a key point being that the derivative is available from the simulated paths. The regularization reduces variance but unlike other techniques such as ridge or lasso regression it does not introduce any bias.

2.1 Simulation of paths

We consider a Black-Scholes model with a single, dividend-free stock, i.e. under the risk-neutral measure Q the stock price dynamics is

$$dS_t = rS_t dt + \sigma S_t dW_t, \quad (1)$$

where both the risk-free interest rate r and the volatility σ are constant. Working in such a sanitized, one-dimensional, log-normal setting is convenient for demonstrating the central concepts of the methods that are to follow. However, it is not a deal-breaker; we are not shoving any rabbits into any hats here – as we will explain along the way.

For any collection of time-points t_i 's, paths of the stock price can be simulated without discretization error

$$S_{t_{i+1}} = S_{t_i} e^{(r - \frac{1}{2}\sigma^2)(t_{i+1} - t_i) + \sigma\sqrt{t_{i+1} - t_i}Z_{i+1}}, \quad (2)$$

where the Z_i 's are independent and $N(0, 1)$. Using the explicit form from equation (2) repeatedly, we get that the derivative of any future value of the process wrt. its initial condition, which we call the path derivative, is

$$\frac{\partial S_{t_j}}{\partial S_{t_i}} = \frac{S_{t_j}}{S_{t_i}} \quad \text{for } i \leq j. \quad (3)$$

2.2 Backward induction and Delta regularization; LSM and Delta LSM

First, to avoid having to go down too many technical rabbit holes, we will in the following assume that our American option is strictly speaking Bermudan, i.e. it can only be exercised on a discrete, finite set of time points $\{t_0 := 0, t_1, \dots, t_{N-1}, t_N := T\}$, an element of which is generically referred to as t_i . However, as we are free to choose both the number of time points and their spacing, we write *American* anywhere it does not confuse.

The backward induction algorithm to value the American option starts at the expiration ($t_N = T$) of the option and works recursively backward in calendar time. For each point on the timeline, we decide whether we exercise the option or hold the option. If we arrive at the expiration date T with a not-yet-exercised option, our choice is trivial: We exercise the option if it is in-the-money, out-of-money options expire worthless.

Standing at a decision point t_i with a not-yet-exercised option, we have to compare what we get from exercising the option now, the intrinsic value, to the value of keeping it alive, the continuation value. Note that our choice is neither "now or never", nor "now or at expiration", but "now or maybe later". Mathematically, we can express the value of the American option, say V^A , via a Bellman-type equation

$$V_{t_i}^A = \max \left\{ \underbrace{h(S_{t_i})}_{=:IV_i}, \underbrace{e^{-r(t_{i+1}-t_i)} \mathbb{E}_{t_i}^Q(V_{t_{i+1}}^A)}_{CV_i} \right\}, \quad (4)$$

where h is the payoff function of the option (so $h(x) = (K - x)^+$ for a put, $h(x) = (x - K)^+$ for a call). The first term inside the curly braces is the intrinsic value, the second term is the continuation value, denoted in shorthand CV_i . Computing the continuation value poses a numerical challenge as it involves the expectation of the option value at a future time conditioned on today's information set. This calculation plays well with backward recursive calculations in a tree, lattice, or grid model. But not so well in the situation with individual simulated stock price paths. This is where the LSM method from Longstaff & Schwartz (2001) comes into play.

By repeated use of equation (4), we can write

$$CV_i = \max_{\tau \in \mathcal{T}_{t_{i+1}; T}} \mathbb{E}_{t_i}^Q(e^{-r(\tau-t_i)} h(S_\tau)), \quad (5)$$

where $\mathcal{T}_{t_{i+1};T}$ denotes the set of stopping times with values in $\{t_{i+1}, t_{i+2}, \dots, T\}$. From eq. (5) we see, we have to calculate the (risk-neutral) expected discounted payoff from exercising optimally in the future. So we can focus on the pay-off from the optimal exercise strategy along the specific path, no reference is made to the V^A -process. By the strong Markov property of S , the conditional expectation on the right-hand side of the equation (5) is a function solely of S_{t_i} , albeit an abstract/general/'very unknown' one. So what we do is approximate it parametrically by a function Φ_i of the form

$$\Phi_i(x) = \sum_{j=0}^M \phi_j(x) \beta_j(i) =: \phi(x) \beta(i), \quad (6)$$

where the ϕ_j 's are some (common across time points, i 's) basis functions that we the modellers choose (and start counting from 0 for later intuitive ease and stack into a row vector for compact notation) and the $\beta_j(i)$'s are parameters, that it is our task to estimate (for each time point). To put a face on the abstract, think of the basis functions as the monomials, i.e. $\phi_j(x) = x^j$. Importantly, Φ_i is linear in parameters – but not in x . Supposing for a minute that we have estimated Φ_i , then for each path (say the l 'th) we know if we should exercise, namely if $h(S_{t_i}^l) > \Phi_i(S_{t_i}^l)$. Working recursively backward, which includes updating the estimated optimal exercise strategy, determines the optimal exercise strategy along each path. The final step – parameter estimation – comes from (for each time point) viewing (for each path; say there are L of them) the discounted realized pay-off from following the (previously estimated) subsequently optimal strategy as the outcome of a random variable whose expected value is the right-hand side of equation (5). This – compactly written – gives us a system of equations

$$\underbrace{e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*})}_{L \times 1} = \underbrace{\phi(S_{t_i})}_{L \times (M+1)} \underbrace{\beta(i)}_{(M+1) \times 1} + \underbrace{\text{noise}}_{L \times 1}, \quad (7)$$

where we write τ_i^* for the L -vector whose l 'th coordinate is the path l , subsequently optimal exercise time as seen from time t_i (i.e. these τ^* 's are updated as we work our way backward). We then estimate the time t_i parameters by minimizing some loss function of

the noise, the sum of squares being a natural choice. So (with $\|x\|^2 = x^\top x$ for a vector x), LSM parameters – or regression coefficients – are estimated as

$$\widehat{\beta}(i) = \arg \min_{\beta(i) \in \mathbb{R}^{M+1}} \|e^{-r\tau_i^*} h(S_{\tau_i^*}) - \phi(S_{t_i})\beta(i)\|^2, \quad (8)$$

i.e. by the standard ordinary least squares formula

$$\widehat{\beta}(i) = (\phi^\top \phi)^{-1} \phi^\top e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*}). \quad (9)$$

2.2.1 Delta regularization

Note that equation (6) tells us that the x -derivative of Φ_i is

$$\frac{\partial}{\partial x} \Phi_i(x) = \sum_{j=0}^M \phi'_j(x) \beta_j(i) =: \phi'(x) \beta(i). \quad (10)$$

This means that in a perfect world – i.e. if the continuation value has the same functional form as Φ_i and we know the parameters – we would also know the derivative of the continuation value wrt. to the underlying – the Delta. It is particularly neat that the expression in equation (10) is also linear in the parameters. However, the approach in the following also works with more complicated Φ_i -functions as long as we have an efficient way of calculating their derivatives (wrt. the underlying). In (dual) neural networks this is exactly what (clever use of) the backpropagation algorithm delivers, which is a key in the differential machine learning approach introduced in Huge & Savine (2020) and Huge & Savine (2021).

Suppose momentarily that we can simulate (at each time point) a vector of random variables, say $Z_S(i)$, whose mean is Delta (eq. (10)). We use this extra information by estimating the regression coefficients β from a regularized loss function

$$Loss = \|e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*}) - \phi(S_{t_i})\beta(i)\|^2 + \lambda \|Z_S(i) - \phi'(S_{t_i})\beta(i)\|^2, \quad (11)$$

where $\lambda = \frac{\| -e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*}) \|^2}{\| Z_S(i) \|^2}$ is chosen such that the two terms in the loss function are of roughly the same size. Minimizing the regularized loss function in equation (11) (see Appendix appendix A) over β , we get the regression coefficients $\widehat{\beta}$,

$$\widehat{\beta}(i) = (\phi^\top \phi + \lambda(\phi')^\top \phi')^{-1} (\phi^\top (e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*})) + \lambda(\phi')^\top Z_S(i)). \quad (12)$$

Comparing equation (12) to equation (9) we see that Delta LSM is an add-on to LSM in a very literal sense.

Of course, the usefulness of the regularization approach just described hinges critically on us being able to produce – i.e. simulate – such mean-Delta Z_S -variables without throwing the baby out with the bathwater, computation time-wise. For this final piece of the puzzle we proceed thus:

$$\begin{aligned} \Delta_i &:= \frac{\partial}{\partial S_{t_i}}(CV_i) = \frac{\partial}{\partial S_{t_i}} \mathbb{E}_{t_i}^Q(e^{-r(\tau_i^* - t_i)} h(S_{\tau_i^*})) \\ &= \mathbb{E}_{t_i}^Q \left(e^{-r(\tau_i^* - t_i)} \frac{\partial}{\partial S_{t_i}} h(S_{\tau_i^*}) \right) = \mathbb{E}_{t_i}^Q \left(e^{-r(\tau_i^* - t_i)} \frac{\partial S_{\tau_i^*}}{\partial S_{t_i}} \frac{\partial h(S_{\tau_i^*})}{\partial S_{\tau_i^*}} \right) \\ &= \mathbb{E}_{t_i}^Q \left(e^{-r(\tau_i^* - t_i)} \frac{\partial S_{\tau_i^*}}{\partial S_{t_i}} h'(S_{\tau_i^*}) \right), \end{aligned}$$

where h' is the derivative of the pay-off function – possibly understood in a generalized (or weak) sense, meaning that it is something that behaves as a derivative should when placed (possibly next to 'something nice') under an integral sign. For call and put options the generalized derivatives are indicator (or Heaviside) functions,

$$h'(x) = \begin{cases} 1_{x > K} & \text{if call} \\ -1_{K > x} & \text{if put.} \end{cases} \quad (13)$$

From equation (3) we get the path derivative in the Black-Scholes model,

$$\frac{\partial S_{\tau_i^*}}{\partial S_{t_i}} = \frac{S_{\tau_i^*}}{S_{t_i}}. \quad (14)$$

This expression allows for one-step simulation and it's specific to the Black-Scholes model. With other, more complicated, processes for the underlying, the idea will (often) also work, but something like equation (14) will have to be applied locally in time along each path, see Giles & Glasserman (2006). Putting all this together, we get our Z_S -variables,

$$Z_S(i) = \frac{S_{\tau_i^*}}{S_{t_i}} h'(S_{\tau_i^*}). \quad (15)$$

Note that Delta LSM the regularization method described above does not add any bias to the estimator but reduces variance, we get a better estimate of the conditional expectation without having to qualify "better" by some bias-variance-trade-off as is usual for regularization methods.

2.3 Forward pass

In the backward induction step, we are using future payoffs as the response variable in the regression to estimate the continuation value. This creates a positive bias (i.e. we overvalue the American option) if we use the backward paths themselves to estimate the American option price; an in-sample or a foresight bias. The larger the sample for the backward induction is, the smaller the bias will be. But for a given sample size (say 10,000 paths) it is far from obvious, what the magnitude of the bias is because of the recursive nature of the backward induction. Effects that are small individually may propagate in nasty ways; studying this is a key part of proving rigorous convergence results about the LSM method, see for instance Stentoft (2004).

However, a simple and computationally cheap way to alleviate this effect is to run an out-of-sample experiment or a forward pass where we simulate a new independent sample of paths, say K of them. For any path, say the k 'th, we march forward in time to decide the optimal exercise point, which is the first time that intrinsic value exceeds the continuation value (and if that does not happen, we get 0 at T), i.e.

$$\hat{\tau}^k = \min\{\min_i \{t_i | h(S_{t_i}^k) > \phi(S_{t_i}^k) \hat{\beta}(i)\}, T\} \quad (16)$$

which is possible because we have estimated the regression coefficients ($\hat{\beta}(i)$) from the in-sample backward induction. Because no matrix inversion is involved, having a large number of paths in the forward pass is computationally much cheaper than having the same number of paths in the backward induction – even without exploiting the embarrassingly parallel nature of the forward pass.

2.4 Valuation

The final step is to average the discounted optimal payoffs from the forward pass to get an estimate of the American option price,

$$\hat{V}_0^A = \frac{1}{K} \sum_{k=1}^K e^{-r\hat{\tau}^k} h(S_{\hat{\tau}^k}^k). \quad (17)$$

The early exercise boundary from the backward induction step gives *some* stopping strategy. The true American option price comes from the stopping strategy that maximizes the discounted expected payoff. Hence, this step will give us a lower bound for the American option price of large (but cheap) K . Because these out-of-sample-paths are independent and the β 's are fixed, we can safely calculate the standard error of such a price estimate with basic statistics.

3 Experimental results

In this section, we compare the performance of the Delta LSM method to that of the standard LSM method for American put option valuation in the Black-Scholes model. A simple example, but still a non-trivial one, as no truly closed-form expression for the put option price is known. As a robustness check, we do this across an ensemble of contract and market parameters. We first demonstrate that our implementation's in-sample LSM results closely match those from the original Longstaff & Schwartz paper – as of course, they should. These benchmark cases use a rather large number of in-sample paths (50,000-100,000). To the naked eye, the differences between LSM and Delta LSM results (both

in- and out-of-sample) for the benchmark cases are small. However, we demonstrate that while yes, benchmark case differences are small enough to ignore in practical applications, the differences are (i) systematically in favour of the Delta LSM method for benchmark in-sample sizes, (ii) dramatically in favour of the Delta LSM for smaller in-sample sizes. More specifically, we find that:

1. $2^{16} = 65,536$ in- and out-sample paths are overkill for both LSM and Delta LSM – but needed for a lazy implementation that uses the same in- and out-of-sample paths.
2. Lowering the number of in-sample paths to $2^{13} = 8,192$ (which seems to be industry standard) – i.e. with 8x lower runtime – both the LSM and Delta LSM give results whose accuracy are acceptable for practical applications.
3. With $2^{10} = 1,024$ in-sample paths – a reduction of 8x paths – the Delta LSM method still gives results that are sufficiently accurate for practical use, but standard LSM method emphatically does not.
4. With $2^{18} = 262,144$ in-sample and out-of-sample paths – Delta LSM exhibits a significantly tighter duality gap, showing a closer-to-optimal exercise strategy even when using 262,144 paths.

3.1 What do you bench?

Table 1 evaluates the Delta LSM against the classical LSM algorithm. Similarly to the original Longstaff and Schwartz paper, we use third-degree monomials ($\phi(x) = (1, x, x^2, x^3)$) and only in-the-money paths for regressions (for out-of-the-money options, the exercise decision – don't! – is trivial). We use multi-dimensional Sobol sequences across paths, meaning that there are technical reasons for preferring numbers of paths that are powers of two. We conduct both in-sample and out-of-sample mean valuation of the American put options in Longstaff & Schwartz (2001) using 100 runs with $2^{16} = 65,536$ in-sample and out-of-sample paths for each run. In each example, we use a strike price $K = 40$ and risk-free rate $r = 6\%$, but as a robustness check we explore various scenarios for

expirations $T \in \{1, 2\}$, volatilities $\sigma \in \{0.2, 0.4\}$, and spot prices $S \in \{36, 38, 40, 42, 44\}$. These are the exact same scenarios as used in Table 1 in the original paper by Longstaff and Schwartz. In Appendix B we demonstrate that our in-sample LSM results closely match those from the original paper.

Table 1: Comparison: LSM and Delta LSM for American Puts

Spot	σ	T	CRR	In-sample		Out-of-sample	
				LSM IS	Delta LSM IS	LSM OFS	Delta LSM OFS
36	0.2	1	4.478	4.476	4.477 (0.005)	4.476 (0.005)	4.477 (0.005)
36	0.2	2	4.840	4.835	4.838 (0.006)	4.835 (0.006)	4.839 (0.007)
36	0.4	1	7.101	7.100	7.101 (0.010)	7.098 (0.010)	7.100 (0.011)
36	0.4	2	8.507	8.504	8.504 (0.013)	8.501 (0.013)	8.504 (0.015)
38	0.2	1	3.250	3.249	3.250 (0.005)	3.248 (0.005)	3.250 (0.006)
38	0.2	2	3.745	3.739	3.741 (0.006)	3.740 (0.006)	3.744 (0.006)
38	0.4	1	6.148	6.147	6.146 (0.010)	6.143 (0.010)	6.146 (0.008)
38	0.4	2	7.668	7.665	7.665 (0.014)	7.664 (0.014)	7.667 (0.014)
40	0.2	1	2.314	2.312	2.313 (0.005)	2.313 (0.005)	2.314 (0.005)
40	0.2	2	2.885	2.881	2.883 (0.006)	2.881 (0.006)	2.884 (0.006)
40	0.4	1	5.312	5.311	5.311 (0.009)	5.308 (0.010)	5.310 (0.010)
40	0.4	2	6.917	6.917	6.917 (0.012)	6.911 (0.011)	6.914 (0.012)
42	0.2	1	1.617	1.616	1.616 (0.004)	1.615 (0.004)	1.616 (0.004)
42	0.2	2	2.213	2.211	2.212 (0.005)	2.209 (0.005)	2.211 (0.005)
42	0.4	1	4.583	4.583	4.582 (0.009)	4.580 (0.009)	4.581 (0.009)
42	0.4	2	6.245	6.242	6.241 (0.013)	6.240 (0.014)	6.242 (0.015)
44	0.2	1	1.110	1.110	1.110 (0.004)	1.108 (0.004)	1.109 (0.004)
44	0.2	2	1.690	1.689	1.690 (0.006)	1.687 (0.006)	1.690 (0.005)
44	0.4	1	3.948	3.947	3.946 (0.010)	3.945 (0.010)	3.947 (0.009)
44	0.4	2	5.642	5.638	5.638 (0.013)	5.638 (0.013)	5.642 (0.013)

Price estimate comparison of American-style put options for in-sample run (step 2: backward induction) and for out-of-sample run (step 3: forward pass) for both LSM and Delta LSM. In this comparison, the parameters are strike price $K = 40$ and risk-free rate $r = 0.06$. The reported numbers for both the in-sample and out-of-sample are the mean and standard error (s.e.) of 100 runs with 2^{16} paths for each run. The CRR column is the benchmark price calculated in a 2,000-steps-per-year Cox-Ross-Rubinstein lattice.

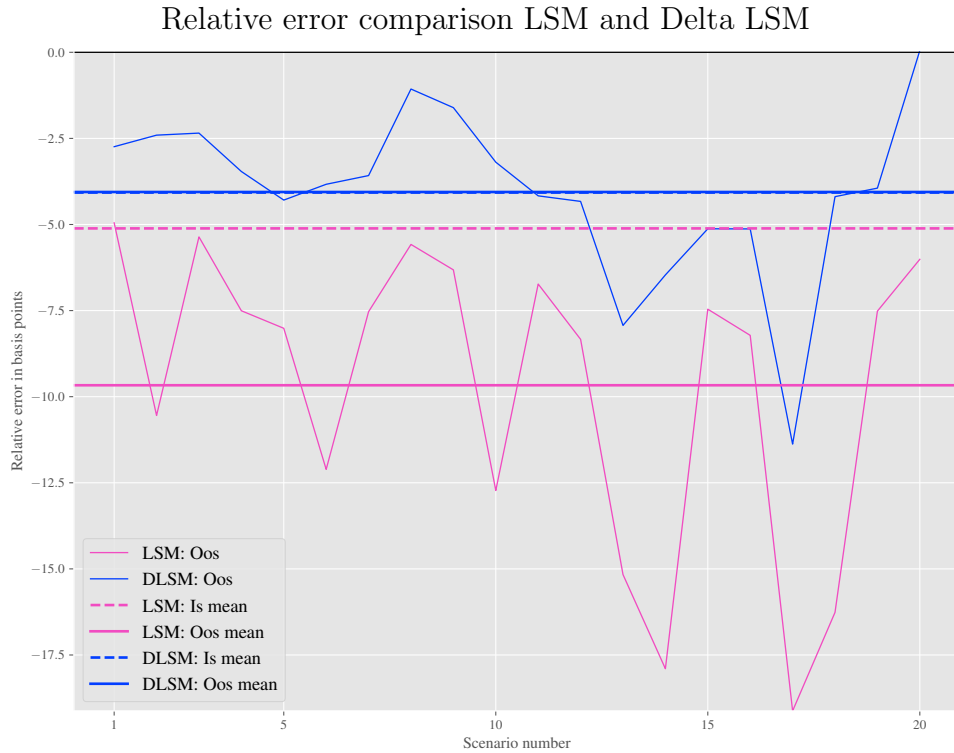


Figure 1: Relative errors in basis points (y-axis) for the price estimates in Table 1. The (categorical) x-axis is the scenario number according to the row number in the table. Dashed is out-of-sample, fully drawn is in-sample. Purple is LSM, blue is Delta LSM. Horizontal lines show averages across scenarios.

We have included results in tabular format to facilitate replication and as an *homage* to the original paper. However, the finer patterns are difficult to see with the naked eye; numbers look pretty similar along each row. So we now zoom in by looking at relative errors – or biases – in basis points,

$$\text{relative error}^{i,j} = 10,000 \frac{\text{price estimate}(\text{method } i, \text{scenario } j) - \text{true price}(\text{scenario } j)}{\text{true price}(\text{scenario } j)},$$

where the method indicator $i \in \{\text{LSM, Delta LSM}\} \otimes \{\text{in-sample, out-of-sample}\}$, and scenario counter $j \in \{1, 2, \dots, 20\}$ refers to the row number in Table 1. Results are presented in Figure 1, which is arguably a bit 'busy', but there is method to the madness:

- On x-axis is the scenario number (row number in Table 1) and on the y-axis is the relative error in basis points.

- Purple is LSM, blue is Delta LSM.
- Dashed curves are in-sample results, and fully drawn curves are out-of-sample results. (For graphical legibility, we connect across the categorical x-variable.)
- Horizontal lines show averages across scenarios, dashed for in-sample, fully drawn for out-of-sample.

The big picture: We deem results to be sufficiently similar across scenarios for us only to comment on their averages, the horizontal lines, in the following. All methods produce errors that are less than 9.7 basis points in absolute value. Typical bid-ask spreads for the options we consider are of magnitude 100 basis points, so all methods produce numbers that are well within range for practical use.

The finer points and patterns: In-sample, Delta LSM and LSM give very similar price estimates (the difference between the two dashed vertical lines is 1.0 basis points). In-sample both methods give negatively biased results, about -5 basis points. The sign of the bias is not apriori clear, as there are conflicting effects in play. By using (only) monomials up to degree three as basis functions, we are restricting our exercise strategies and losing flexibility regarding the early exercise boundary – albeit in a complicated way. And since the true price comes exactly from the exercise strategy that maximizes the expected discounted payoff from following that strategy, we have a negative source of bias. However, a source of positive bias is that in-sample we use the *exact same* paths to evaluate the performance of our exercise strategy as we used to determine it; the smaller the sample size, the more pronounced the bias will be. The results show that with $\sim 50,000$ paths, the overall effect is negative. (One cannot simply, yet, conclude, that the in-sample bias is negligible; the small negative bias might come as a sum of a large positive and a 'large' negative number.) Let us now turn to the out-of-sample results. We know that this should be negatively biased – which is indeed what we see. But now there is a visibly clear difference between LSM and Delta; LSM has a bias of -9.7 basis points, Delta LSM only -4.1 basis points. Any in-sample estimation noise/inaccuracy will be into negative bias out-of-sample. This shows that the Delta LSM gives more accurate

estimates of the optimal exercise strategy. Notice further the difference between in-sample and out-of-sample results is minuscule for Delta LSM (the dashed and fully drawn blue horizontal lines almost coincide; 0.02 basis point difference) but not for LSM (4.55 basis points difference between the purple lines). So not only does Delta LSM give us a better strategy out-of-sample, but it is also trustworthy price-wise in-sample, whereas LSM is more unrealistically optimistic about its performance.

3.2 Duel: Size matters

Using the same experimental design as in the previous section, we now investigate what happens when we decrease the in-sample sample size; first from 2^{16} down by a factor of 8 to $2^{13} = 8,192$ ($\sim 10,000$), which is commonly used in practice, and then down by a further factor 8 to $2^{10} = 1,024$ ($\sim 1,000$), a nice, (almost) round and intuitively large number.

The results are shown in Figure 2. The left-hand panel is Figure 1 reproduced for comparison. Looking at the 8,192-paths middle panel, we see that all methods still give practically usable results (absolute biases less than 22 basis points). But while the performance of Delta LSM is near-as-makes-no-difference unchanged from the earlier 65,536-paths case, the LSM is starting to show small-sample shortcomings. Increased positive small-sample bias leads to a price estimate closer to the true one, but this is a result of two wrongs making a right; using LSM the estimated exercise strategy performs considerably worse than that from the Delta LSM. Turning to the 1,024-paths right-hand panel we see that the positive small-sample bias makes the Delta LSM visibly optimistic about its performance (bias +14 basis points), but that using its estimated exercise strategy still gives acceptable performance (out-sample-bias -36 basis points). However, for LSM, things are not well. In-sample it has a bias of +135 basis points, while the strategy's out-of-sample performance is -89 basis points. To see why this is dangerous consider this simple example: The true price of the option is 10. You can buy it at 10.05 and sell it at 9.95 – a bid-ask spread of 100 basis points. Using your in-sample LSM estimate, you reckon that the true value of the option is 10.135. So you buy the option at 10.05, short the replicating




Figure 2: Relative pricing errors across scenarios for different in-sample sample sizes. Each panel corresponds to a different size, indicated on its top. On the y-axes are relative pricing errors (biases) in basis points, on the x-axes are scenario numbers (the initially deep out-of-the-money scenarios 17-20 have been excluded from the graphs because relative errors are very sensitive to the treatment of lacking in-the-money paths). Purple is for LSM, blue is for Delta LSM. Dashed curves show in-sample performance, fully drawn curves show out-of-sample performance (estimated from 65,536 paths). Horizontal lines indicate averages across scenarios with the same dashed/full convention.

portfolio for your estimated optimal strategy, and follow said strategy for exercising the option that you hold – all in the firm belief that that will give you an arbitrage profit of $10.135 - 10.05 = 0.085$. However, the -89 basis points out-of-sample performance tells us that the expected discounted payoff from following your estimated optimal strategy is only 9.91 . So rather than making 0.085 for sure, you will lose 0.14 on average (and there will even be some randomness to your loss).

3.3 Dual: Tightly bound

Even with an infinite number of out-of-sample paths, the LSM and the Delta LSM methods will only produce a lower bound on the American option price. Whether or not that is close to the true price depends on the problem at hand, and on the choice (number and functional form) of basis functions. Suppose that for some reason safety is our main concern; it is really important for us that we have both upper and lower bounds for the American option price. The primal-dual algorithm from Andersen & Broadie (2004) provides a method to turn a lower bound (and its associated exercise strategy) – be that from LSM or Delta LSM – into an upper bound on the arbitrage-free price of the American option. The primal-dual algorithm goes in three steps:

1. Estimate an optimal exercise strategy $\hat{\tau}$ using a primal algorithm such as Delta LSM or LSM, which gives a lower bound of the true price

$$\hat{V}_0^A = \frac{1}{K} \sum_{k=1}^K e^{-r\hat{\tau}^k} h(S_{\hat{\tau}^k}^k) \leq \sup_{\tau \in \mathcal{T}_{0,T}} \mathbb{E}^Q(D_{0,\tau} h_\tau), \quad (18)$$

where for compactness' sake we set $D_{t,u} = e^{-r(u-t)}$.

2. Use the exercise strategy, to derive an upper bound by nested simulation:
 - Simulate K_{outer} paths; for each path k' do:
 - If $\hat{\tau}^{k'}(t_i)$ instructs us to exercise, then set $\hat{\tau}^{k'}(t_i) = t_i$ and do the following:
 - * Set $\hat{V}'_{t_i} = D_{0,t_i} h_{t_i}$
 - * Simulate K_{nested} paths (starting from the situation at t_i and denote each path $k^* = 1, 2, \dots, K_{nested}$) to estimate the value instructed by the optimal stopping time $\hat{\tau}_{i+1}$ from t_{i+1} to expiration T :

$$\mathbb{E}_{t_i}^Q(\hat{V}'_{t_{i+1}}) \approx \frac{1}{K_{nested}} \sum_{k^*=1}^{K_{nested}} D_{0,\hat{\tau}_{i+1}} h(S_{\hat{\tau}_{i+1}}^{k^*}), \quad (19)$$

where to streamline notation, we suppress the k dependence on $\hat{\tau}_{i+1}$.

- If $\hat{\tau}^{k'}(t_i)$ instruct continuation, then $\hat{\tau}^{k'}(t_i) > t_i$ and do the following

* Simulate K_{nested} paths to estimate:

$$\hat{V}'_{t_i} = \mathbb{E}_{t_i}^Q(\hat{V}'_{t_{i+1}}) \approx \frac{1}{K_{nested}} \sum_{k^*=1}^{K_{nested}} D_{0, \hat{\tau}_{t_{i+1}}} h(S_{\hat{\tau}_{t_{i+1}}}^{k^*}) \quad (20)$$

– Build the martingale strategy:

$$\widehat{M}'_{t_{i+1}} = \widehat{M}'_{t_i} + \widehat{V}'_{t_{i+1}} - \mathbb{E}_{t_i}^Q[\widehat{V}'_{t_{i+1}}], \quad \text{with } \widehat{M}'_{t_0} = 0.$$

The martingale \widehat{M}' is our estimate of the optimal martingale of path k' out of the set of right-continuous martingales $\mathcal{M}_{t,0}$ starting at zero².

3. Averaging over our estimates of the optimal martingale for each outer path k' to construct the upper bound of the American option:

$$\sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E}_t^Q[D_{t,\tau} h(S_\tau)] \leq \inf_{M \in \mathcal{M}_{t,0}} \mathbb{E}_t^Q[\sup_{t \leq u \leq T} (D_{t,u} h(S_u) - M_u)] \quad (21)$$

$$\approx \frac{1}{K_{outer}} \sum_{k'=1}^{K_{outer}} \max_{u \in \{t,T\}} (D_{t,u} h(S_u^{k'}) - \widehat{M}'_u). \quad (22)$$

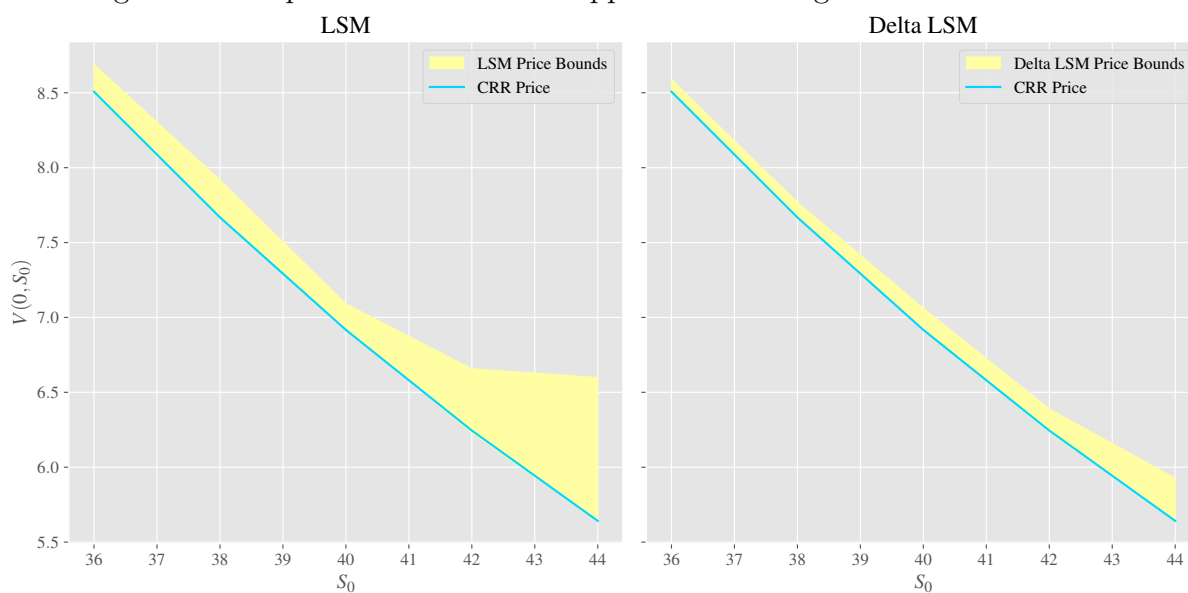
We use the dual algorithm to construct both lower and upper bound for American put options with expiry $T = 2$, strike $K = 40$, volatility $\sigma = 0.4$, and risk-free rate $r = 0.06$ across spot-prices $S \in \{36, 38, 40, 42, 44\}$. In Figure 3, the duality gap between the lower and upper bound is shaded yellow, and the benchmark price is cyan. (We provide the full table of lower and upper bound prices across all the scenarios in Table 2.)

To construct the lower and upper bound, we use 2^{18} Sobol paths for the primal algorithm (LSM or Delta LSM), 2^{11} outer Sobol paths, and 2000 nested non-Sobol paths.

Across all spot prices, we see a tighter duality gap for Delta LSM in Figure 3, as a result of a closer-to-optimal exercise strategy using Delta LSM, even with the use of $2^{18} = 262,144$ in-sample and out-of-sample paths. The duality gap sheds light on the benefit of using

²The construction comes from the Doob-Meyer decomposition, however, let's not get stuck on mathematical formality

Figure 3: Comparison of lower and upper bound using LSM and Delta LSM



Price estimate comparison of American-style put options using the primal-dual algorithm for both LSM and Delta LSM using a cubic polynomial as basis. In this comparison, the parameters are strike price $K = 40$, the expiry $T = 2$, the volatility $\sigma = 0.4$, and risk-free rate $r = 0.06$. To generate the numbers, we run 2^{18} paths for the primal algorithm, 2^{11} for the outer paths in the dual algorithm, and 2^{11} inner paths in the dual algorithm. The CRR is the benchmark price and the lines are linearly interpolated.

Delta LSM on bigger sample sizes, challenging the notion of marginal benefits as suggested in sections 3.1 and 3.2. In the absence of a convergent lattice method, Delta LSM stands out by providing a markedly improved estimate of the range of true value by giving a tighter no-arbitrage bound.

4 Conclusion

In this paper we have described and tested (with favorable results) the Delta LSM method which enhances the classical LSM method by use of the simulated derivative of the continuation value. While this is not the be-all – and certainly not the end-all – of simulation-based American option pricing, it is a cheap-in-several-ways method that can easily – and without any downside risk that we are aware of – be added on to existing domain-specific pricing methods and libraries.

References

- Andersen, L. & Broadie, M. (2004), ‘Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options’, *Management Science* **50**(9), 1222–1234.
- Andersen, L., Lake, M. & Offengenden, D. (2016), ‘High-performance American option pricing’, *Journal of Computational Finance* **20**(1), 39–87.
- Becker, S., Cheridito, P., Jentzen, A. & Welte, T. (2021), ‘Solving high-dimensional optimal stopping problems using deep learning’, *European Journal of Applied Mathematics* **32**(3), 470–514.
- Frandsen, M. G., Pedersen, T. C. & Poulsen, R. (2022), ‘Delta force: Option pricing with differential machine learning’, *Digital Finance* **4**, 1–15.
- Giles, M. & Glasserman, P. (2006), ‘Smoking adjoints: fast Monte Carlo greeks’, *Risk* (January), 88–82.

Table 2: Tightness of bounds using LSM and Delta LSM for American Puts

Spot	σ	T	CRR	LSM Primal	LSM Dual	Delta LSM Primal	Delta LSM Dual
36	0.2	1	4.478	4.475	4.592	4.476	4.520
36	0.2	2	4.840	4.836	4.996	4.840	4.887
36	0.4	1	7.101	7.101	7.235	7.099	7.178
36	0.4	2	8.507	8.499	8.690	8.501	8.589
38	0.2	1	3.250	3.248	3.358	3.249	3.289
38	0.2	2	3.745	3.740	3.879	3.743	3.795
38	0.4	1	6.148	6.150	6.292	6.152	6.227
38	0.4	2	7.668	7.669	7.919	7.673	7.770
40	0.2	1	2.314	2.312	2.417	2.313	2.351
40	0.2	2	2.885	2.885	3.014	2.888	2.933
40	0.4	1	5.312	5.314	5.450	5.312	5.381
40	0.4	2	6.917	6.926	7.091	6.928	7.062
42	0.2	1	1.617	1.614	1.735	1.616	1.650
42	0.2	2	2.213	2.210	2.362	2.212	2.257
42	0.4	1	4.583	4.577	4.697	4.579	4.640
42	0.4	2	6.245	6.231	6.658	6.235	6.390
44	0.2	1	1.110	1.110	1.981	1.110	2.143
44	0.2	2	1.690	1.687	3.165	1.690	3.135
44	0.4	1	3.948	3.947	4.180	3.952	4.023
44	0.4	2	5.642	5.641	6.599	5.642	5.927

Price estimate comparison of American-style put options using the primal-dual algorithm for both LSM and Delta LSM using a cubic polynomial as basis. In this comparison, the parameters are strike price $K = 40$ and risk-free rate $r = 0.06$. To generate the numbers, we run 2^{18} Sobol paths for the primal algorithm, 2^{11} for the outer Sobol paths in the dual algorithm, and 2^{11} inner non-Sobol paths in the dual algorithm. The CRR column is the benchmark price.

Huge, B. & Savine, A. (2020), ‘Differential machine learning: the shape of things to come’, *Risk* (September).

Huge, B. & Savine, A. (2021), ‘Axes that matter: PCA with a difference’, *Risk* (September).

Lind, P. (2022), Overcoming the Feature Selection Issue in the Pricing of American Options, working paper, SSRN: Social Science Research Network.

Longstaff, F. A. & Schwartz, E. S. (2001), ‘Valuing American Options by Simulation: A Simple Least-Squares Approach’, *Review of Financial Studies* **14**(1), 113–147.

Stentoft, L. (2004), ‘Convergence of the least squares Monte Carlo approach to American option valuation’, *Management Science* **50**(9), 1193–1203.

A Optimal Regression Coefficients for Delta LSM

To solve a regularized least squares problem of the form

$$\arg \min_{\beta} \left(\|Y - \phi\beta\|^2 + \sum_j \lambda_j \|Z_j - \phi_j\beta\|^2 \right), \quad (23)$$

we differentiate the expression inside the large parentheses and put it equal to zero

$$\begin{aligned} & \frac{\partial}{\partial \beta} (Y^\top Y - Y^\top \phi\beta - (\phi\beta)^\top Y + (\phi\beta)^\top (\phi\beta)) \\ & + 2 \sum_j \lambda_j (Z_j^\top Z_j - Z_j^\top \phi_j\beta - (\phi_j\beta)^\top Z_j + (\phi_j\beta)^\top (\phi_j\beta)) \\ & = -\phi^\top Y + \phi^\top \phi\beta + \sum_j \lambda_j (-\phi_j^\top Z_j + \phi_j^\top \phi_j\beta) = 0, \end{aligned}$$

so that

$$\hat{\beta} = \left(\phi^\top \phi + \sum_j \lambda_j \phi_j^\top \phi_j \right)^{-1} (\phi^\top Y + \sum_j \lambda_j \phi_j^\top Z_j). \quad (24)$$

B Comparison of Tables 1

As a sanity check, in Figure 4 we compare the results in our Table 1 to those reported in Table 1 in the original Longstaff & Schwartz-paper. The short version: We are quite happy with the results.

The longer version: European put option prices calculated from the Black-Scholes formula (which we don't actually report in our Table 1) are spot on (top left panel). An absolute minimal requirement. For the 'true' 50-dates Bermudan option prices (top right panel), which we calculate with a Cox-Ross-Rubinstein binomial model and Longstaff and Schwartz by a finite difference method, there are some last (of three) decimal discrepancies. We have no definite explanation for this, but deem it to be unimportant in magnitude. It should be noted that the difference between the 50-dates-per year Bermudan option price and 'truly continuously exercisable' American option price is clearly visible (top right panel) and of non-negligible magnitude (bottom left panel) – which is why we compare to (our) Bermudan prices in subsequent analysis. The bottom left panel shows that there is some variation in reported in-sample LSM price estimates, but nothing that appears systematic or more pronounced than what we would expect from the estimated standard errors. Finally, the bottom right panel shows the ratio of reported in-sample standard errors for LSM price estimates. Some variation here is inevitable because for each parameter combination, the result reported by L&S is based on a single run with 10^5 paths. Our results are based on 100 runs each with $2^{16} = 65,536$ paths (there are Sobol simulation reasons for us using a power of two). But we would expect the average of standard errors ratios to be $\sqrt{2^{16}/10^5} = 0.81$ – which near-as-makes-no-difference is what we see (the black dashed line vs. the fully drawn blue one).

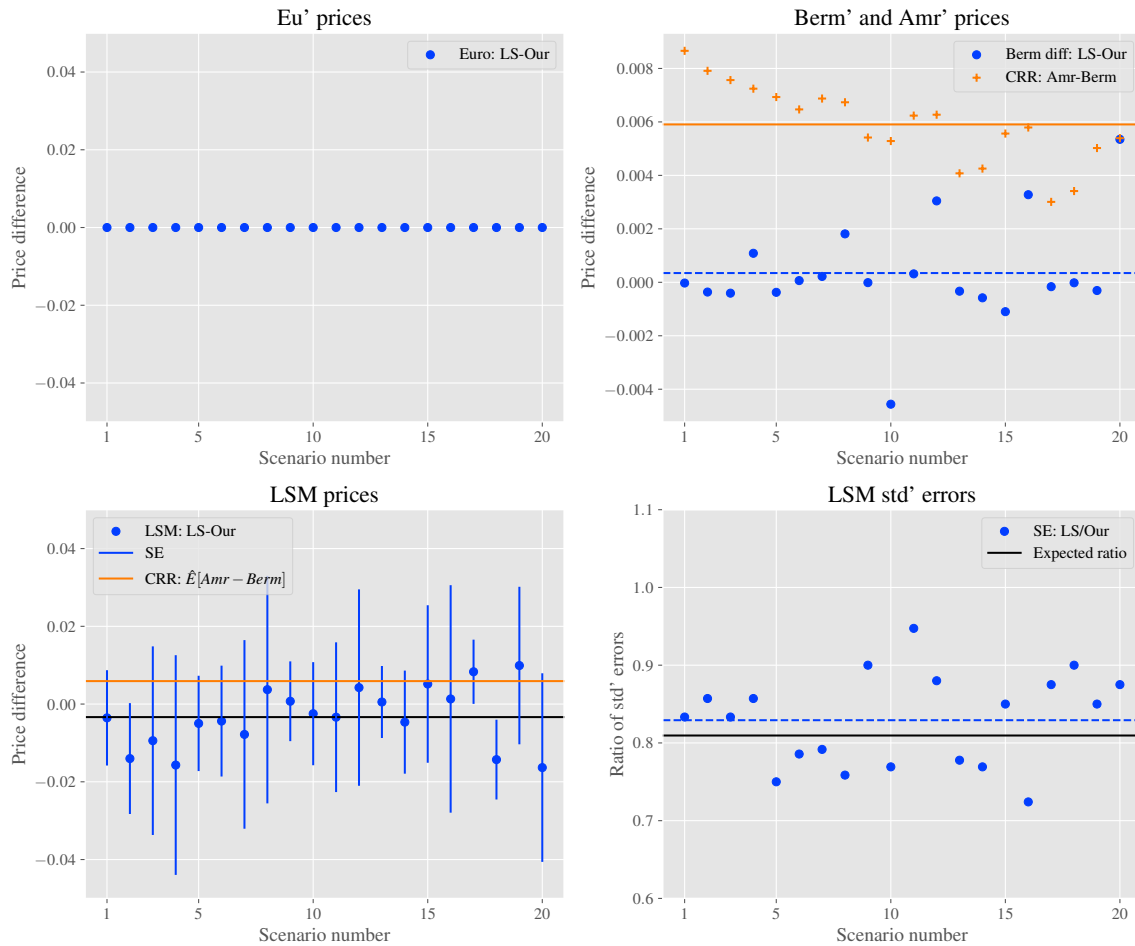


Figure 4: Comparison of our Table 1 results to those reported in Table 1 in the original Longstaff & Schwartz-paper. The experiment numbers on the x-axes correspond to the parameter combinations in Table 1 indexed by row number. The top left panel are differences in European options prices. In the top right panel, the blue circles are differences between L&S's reported 50-dates-per-year Bermudan finite difference prices and our calculations from prices from a 2000-steps-per-year CRR model. The orange pluses are differences between CRR Bermudan and American options prices. Dashed horizontal lines are averages over the 20 experiments. In the bottom left panel, the circles show differences between reported LSM in-sample estimated option prices. The vertical lines are \pm L&S's reported standard error and the dashed orange horizontal line is the one from from the top right panel. In the bottom right panel, the circles are the ratios of reported in-sample standard errors, with the dashed black line showing their average, and the blue line indicating what we'd expect this to be, $\sqrt{2^{16}/10^5} = 0.81$.