Discretionary Exercise by Conditional-Density Reweighting of Monte Carlo Paths

Gregory Pelts

1 November 2026

Abstract

We study a Monte Carlo procedure for discretionary exercise that reweights simulated futures with short-step transition densities. Its principal advantage over continuation-value regression is a general, model-independent exercise interface: the model supplies conditional weights, the contract supplies cashflows, and the exercise engine requires no choice of regression variables or basis functions tied to internal model factors. Prices remain dependent on the specified dynamics. The construction belongs to the stochastic-mesh family and preserves each original path's selected cashflows for final valuation. We develop the conditional-density identity and its numerical implementation, illustrate a surrenderable Asian contract, and compare Black–Scholes and Heston exercise prices with tree and finite-difference references.

Keywords: American options; Bermudan options; discretionary exercise; Monte Carlo; stochastic mesh; conditional density; least-squares Monte Carlo.

AI-use disclosure. OpenAI tools assisted manuscript drafting, editing, reference checking and development of the research software.

1 Problem and relation to existing methods

The practical objective is a common exercise procedure for different products and market models. An exercise feature should not require the product author to select a new collection of regressors whenever the simulator introduces a volatility factor, changes its state representation or adds another source of risk. This motivates the separation used here: the model provides conditional weights, the contract provides possible payments, and a shared exercise engine compares the available actions.

Under a pricing measure $\mathbb{Q}$, let X_t be a Markov market state, augmented where necessary by contract information sufficient to determine future cashflows. A holder may exercise at dates $0 < t_1 < \cdots < t_M = T$. For clarity, discounting is deterministic at constant rate r. If exercise pays $g_k(X_{t_k})$ and terminates the contract, the discrete stopping problem is

$$V_M(x) = g_M(x), \tag{1}$$

$$C_k(x) = \mathbb{E}[e^{-r\Delta_k} V_{k+1}(X_{t_{k+1}}) \mid X_{t_k} = x], \qquad \Delta_k = t_{k+1} - t_k, \tag{2}$$

$$V_k(x) = \max\{g_k(x), C_k(x)\}. \tag{3}$$

Receiving a surrender payment instead of retaining a claim is one instance of the same comparison. Multiple contractual alternatives require their own conditional cashflow values and state updates; the executable integration considered below is narrower than that general problem.

1.1 Comparison with continuation-value regression

Longstaff and Schwartz [2] estimate continuation from simulated future cashflows by regressing on functions of the current state. Tsitsiklis and Van Roy [3] develop regression-based approximate value iteration in a chosen family of value functions. Both approaches accommodate

multiple factors and path dependence; they are not confined to a single stock variable or vanilla payoffs. Their practical specification nevertheless includes the variables and functions used to represent continuation. In stochastic volatility, for example, a basis in stock alone need not capture the influence of latent variance; contract history may add further relevant inputs.

Here the conditional averaging rule comes from the model's transition law. A short transition connects a target state to the already simulated futures of other paths. Reweighting their cashflows produces continuation estimates without fitting a global functional form for continuation. The same operation can be applied to another payoff without selecting new powers, interactions or other continuation regressors. A direct nested Monte Carlo method also avoids a continuation basis, but launches fresh conditional futures at each target and decision date; the present construction reuses a common path bank.

Table 1: The implementation distinction from continuation-value regression. Both approaches still require sufficient information for conditional pricing.

Task	Continuation-value regression	Conditional-density path reuse
Estimate continuation	Fit cashflow values in selected state features and approximation functions.	Average donor cashflows with model-supplied conditional weights.
Change the model	Reassess whether the chosen features and function family capture the new conditional value.	Supply the new transition and sampling-density calculation; retain the exercise logic.
Change the contract	Specify cashflows and sufficient history; reassess the continuation representation.	Specify cashflows and sufficient history; retain the conditional averaging rule.
Assess numerical error	Assess approximation-family error and statistical estimation error.	Assess density approximations, donor support and statistical estimation error.

1.2 Generality and the model interface

Model independence here concerns the exercise algorithm and its interface. The product author does not need to select functions of the simulator's internal factors to build a continuation regression. The model provider still needs a sufficient state and a compatible conditional transition law; changing the model changes the conditional weights and can change exercise decisions and prices. Transition-density evaluation or a controlled approximation is therefore a substantive requirement, rather than something furnished by every path simulator automatically.

Contract dependence is likewise confined to the appropriate interface, not eliminated. The engine must compare the actual contractual actions and preserve the target's relevant history. Section 6.2 makes this concrete for a surrenderable Asian payoff. Generality of the conditional construction does not imply that every history-dependent product is already supported by the current integration.

The optional moment controls used below are distinct from a basis defining the continuation approximation. The uncontrolled conditional estimator is already specified by the density weights. A provider may additionally supply variables with known conditional means to reduce its error. The Heston example uses such variables and estimates local control coefficients by regression; it does not require a product-specific regression representation of the value function. Removing a mandatory continuation basis is the central advantage, not a claim that regression calculations are prohibited or that density estimation is universally cheaper.

1.3 Relation to stochastic mesh

Density ratios of this kind belong to the stochastic-mesh family. Broadie and Glasserman [1] formulate transition-density divided by sampling-density weights, including marginal and empirical-mixture sampling densities, and discuss control variates and policy evaluation. Neither density reweighting nor inner controls are claimed here as new principles. The focus is a particular short-step implementation supporting that reusable exercise interface: localized conditional estimation, individual exercise decisions and preservation of each original path's realized cashflows. The model-specific work lies in obtaining suitable density weights and, optionally, conditional moments; the contractual action comparison remains common.

The development proceeds from the density identity to normalized conditional estimation, individual exercise decisions and optional moment controls. An exact conditional-call calculation and a surrenderable Asian contract illustrate the construction. Numerical examples then compare Black–Scholes and Heston exercise prices with tree and finite-difference references, with diagnostics for density approximations, donor support and computational cost. This distinguishes the general exercise interface from the numerical approximations used by a particular model provider.

2 Conditioning through the first subsequent simulation step

2.1 The marginal-density identity

Fix an exercise time t and a small simulation step $h > 0$ with $t+h \le t_{k+1}$. A donor path supplies $Y = X_{t+h}$ and a later cashflow R, expressed in units of value at t_{k+1}. For this section the future policy generating R is fixed. Write $p_h(y \mid x)$ for the transition density from the target x, and $q_{t+h}(y)$ for the marginal density of the donor endpoint generated from the valuation-date initial condition. Densities must use the same reference measure. Markov sufficiency means that, after fixing any target contractual state, the donor's future conditional on Y has the law required for that target. Let $m(y) = \mathbb{E}[R \mid Y = y]$.

Proposition 1 (Conditional path reuse)**.** *Suppose the target transition is absolutely continuous with respect to the donor marginal on its full support, and the weighted cashflow is integrable. Then*

$$\mathbb{E}[R \mid X_t = x] = \mathbb{E}\left[\frac{p_h(Y \mid x)}{q_{t+h}(Y)} R\right]. \tag{4}$$

Proof. Conditioning first on Y gives

$$\mathbb{E}\left[\frac{p_h(Y \mid x)}{q_{t+h}(Y)} R\right] = \int q_{t+h}(y) \frac{p_h(y \mid x)}{q_{t+h}(y)} m(y)\, dy = \int p_h(y \mid x) m(y)\, dy.$$

The last expression is the target conditional expectation by the Markov property. Conditioning at a continuously distributed state is understood through its transition kernel, not as an event that must occur in the sample. □

The denominator in (4) is the density of the *sampled endpoint*. It is not $p_h(Y_j \mid X_t^j)$, the transition density conditional on donor j's own preceding state. In a one-dimensional stock model this is the spot marginal at $t + h$; in a stochastic-volatility model the ideal denominator is the joint marginal of the relevant market state.

For independent donors and a fixed target, the unnormalized average of the right-hand side is unbiased under these assumptions. This statement does not extend automatically to a normalized ratio, an approximate marginal, a truncated donor set, or a policy fitted recursively on the same paths.

2.2 Normalized rows and original path identities

For target path i, form weights

$$w_{ij} = \frac{p_h(Y_j \mid X_t^i)}{q_{t+h}(Y_j)}, \qquad \alpha_{ij} = \frac{w_{ij}}{\sum_{\ell \neq i} w_{i\ell}}, \qquad j \neq i. \tag{5}$$

The conditional estimate without controls is

$$\widehat{C}_i = e^{-r\Delta_k} \sum_{j \neq i} \alpha_{ij} R_j. \tag{6}$$

The same original path ID is excluded even if observations are reordered for searching or computation. Exclusion prevents the most direct use of the target's own subsequent realization in its estimate.

Row normalization preserves constant responses exactly, including after deterministic donor truncation, provided the retained row has positive mass. With exact densities, a fixed target, independent donors, fixed h, a fixed future policy and no donor truncation, the law of large numbers gives numerator convergence to the desired conditional value and denominator convergence to one. Thus the normalized estimator is consistent in that setting, but is generally biased at finite sample size. A fixed truncation instead targets the corresponding normalized restricted integral. If a screened row is empty, the implementation recomputes it with all nonself donors and stable log-weight normalization.

Two time intervals have different roles. The transition density spans the first simulation step h after the decision. Cashflows in (6) are discounted over the full exercise interval Δ_k. Increasing the number of simulation steps does not add exercise dates unless the contract schedule is also changed.

2.3 Contract histories

For a path-dependent contract, a donor's already accumulated contract history need not equal the target's history. It cannot be substituted indiscriminately. If a target history a_i can be held fixed while the donor future is replayed, the response becomes $R_{ij} = R(a_i, \text{future}_j)$ and the same density identity applies to that target-specific functional. Alternatively, a sufficient contract state must be included in the conditioning state with a compatible transition law. Section 6.2 illustrates the first construction at a single decision date. General history-dependent branching is not implemented by the current call/put integration.

3 Backward decisions and realized cashflow valuation

The recursion stores $B_{k,i}$, the selected realized cashflow on original path i, valued at t_k. It does not store a conditional mean as that path's payment. Initialize $B_{M,i} = g_M(X_{t_M}^i)$. Going backward, estimate continuation from the donor responses $B_{k+1,j}$, then set

$$I_{k,i} = \mathbf{1}\{g_k(X_{t_k}^i) > 0, \quad g_k(X_{t_k}^i) \geq \widehat{C}_{k,i}\}, \tag{7}$$

$$B_{k,i} = I_{k,i} g_k(X_{t_k}^i) + (1 - I_{k,i}) e^{-r\Delta_k} B_{k+1,i}. \tag{8}$$

The positive-payoff condition in (7) is the current call/put convention; general actions need their contractual comparison. The price is

$$\widehat{V}_0 = \frac{1}{N} \sum_{i=1}^{N} e^{-rt_1} B_{1,i}. \tag{9}$$

Every path has an individual exercise decision. The full original donor population remains available during backward induction, including paths that would be inactive under an earlier exercise decision during forward replay.

For a single decision, the realized exercise increment on a path is $I_i\{g(X_t^i) - e^{-r\Delta}B_{\text{future},i}\}$. The final mean of these increments is an accounting decomposition of the price. The realized future difference itself is not observable at the decision date. Using its sign to choose exercise, or to automatically classify targets outside a computational bucket, would use future information. Computational sparsity therefore does not replace the conditional comparison.

The role of conditional accuracy is particularly clear for one decision. If the true continuation is $C(x)$, a wrong choice induced by $\widehat{C}(x)$ loses

$$|g(x) - C(x)|\mathbf{1}\{\text{wrong choice}\} \le |\widehat{C}(x) - C(x)|. \tag{10}$$

The inequality follows because an incorrect comparison places $g(x)$ between $C(x)$ and $\widehat{C}(x)$. It is not an error bound for the full implemented recursion without further control of future-policy errors.

3.1 Algorithm

1. Simulate a common path bank on the chosen physical grid. Save the market state at each exercise date and at its first subsequent simulation step, together with states needed for future contractual payments.
2. Set each terminal response to that path's terminal payoff.
3. At each earlier exercise date, form nonself donor weights in a common coordinate system, apply the chosen density approximation and screening, and normalize each target row.
4. Estimate each target's continuation, optionally with the conditional-moment controls below. Compare the estimate with that target's immediate exercise value.
5. Retain the target's own exercise payment or its own discounted future response according to the comparison. Record its decision.
6. Average the original paths' discounted selected payments. In a contractual runtime, replay the recorded decisions and value the resulting payment records.

3.2 Shared-sample interpretation

Direct self-exclusion is not complete statistical separation. A donor's later decision may itself have been learned from the target path through other rows and dates. Consequently, the implementation is a shared-bank policy estimator. Its final cashflows are not independent observations of a fixed, externally fitted policy. In particular, a reported $s(B)/\sqrt{N}$ is a descriptive cashflow statistic, not an independently estimated standard error of the complete fitting-and-pricing procedure. The same-bank price is not claimed to be a rigorous lower or upper bound on the optimal price.

4 Marginal approximations, localization and controls

4.1 What replacing the marginal changes

Suppose the actual donor density is q, but the weights divide by a positive approximation $\widetilde{q}$. Holding the transition p_h and future policy fixed, set $a(y) = q(y)/\widetilde{q}(y)$. The large-sample normalized expression is

$$\widetilde{m}(x) = \frac{\int p_h(y \mid x)a(y)m(y)\,dy}{\int p_h(y \mid x)a(y)\,dy}. \tag{11}$$

Its difference from $\mathbb{E}_{p_h(\cdot|x)}[m(Y)]$ is exactly

$$\frac{\mathrm{Cov}_{p_h(\cdot|x)}(m(Y), a(Y))}{\mathbb{E}_{p_h(\cdot|x)}[a(Y)]}. \tag{12}$$

Thus a multiplicative constant common to a row cancels; a spatially varying factor generally does not. An approximately constant ratio over the transition neighborhood motivates a local approximation. It does not make that approximation an exact change of measure at a fixed positive step.

A very short transition concentrates near its starting state. Away from singular boundaries, a regular d-dimensional diffusion has a typical local volume of order $h^{d/2}$. The resulting heuristic local donor count is of order $Nh^{d/2}$ times the local marginal density and diffusion-volume factor. Reducing h can improve a local density approximation while worsening donor support. Refining simulation time alone therefore does not establish convergence. This scaling is an interior heuristic, not a theorem for a degenerate variance boundary or for the complete algorithm.

The coordinate convention matters. A constant marginal with respect to dS is not constant with respect to $d\log S$. Exact numerator and denominator densities transform together; replacing only the denominator by a constant breaks that invariance. The Heston calculation below uses physical (S, v) coordinates.

4.2 Optional conditional-moment controls

The density-weighted estimate in (6) requires no regression basis. This subsection adds variance-reduction controls whose conditional means are supplied by the model. Choosing these optional controls is separate from choosing a function family to represent a contract's continuation value.

Let $\phi(Y) \in \mathbb{R}^p$ have a known conditional mean $\mu_i = \mathbb{E}[\phi(Y) \mid X_t = X_t^i]$, and define $z_{ij} = \phi(Y_j) - \mu_i$. With normalized retained weights, write

$$\bar{B}_i = \sum_j \alpha_{ij} B_{k+1,j}, \qquad \bar{z}_i = \sum_j \alpha_{ij} z_{ij}, \tag{13}$$

$$\Sigma_i = \sum_j \alpha_{ij} (z_{ij} - \bar{z}_i)(z_{ij} - \bar{z}_i)^\top, \qquad c_i = \sum_j \alpha_{ij} (z_{ij} - \bar{z}_i)(B_{k+1,j} - \bar{B}_i). \tag{14}$$

For full rank, $\widehat{\beta}_i = \Sigma_i^{-1} c_i$ gives

$$\widehat{C}_{k,i} = e^{-r\Delta_k}(\bar{B}_i - \widehat{\beta}_i^\top \bar{z}_i). \tag{15}$$

This is the weighted regression intercept evaluated at the known conditional control mean. It uses the model's analytic expectation, not empirical recentering of simulated stock paths. In Black–Scholes the control is $S' - S_i e^{(r-q)h}$. In the Heston example the controls are S', v' and $S'v'$, each centered at its analytic mean under the simulated short-step law.

The implementation standardizes the covariance to a correlation matrix and discards directions unresolved at floating-point precision. At deficient rank it uses the corresponding minimum-norm slope in standardized coordinates; with rank zero it retains the weighted mean. There is no ridge penalty, coefficient cap or continuation-value clipping. Sparse rows may leave extrapolation to the target mean poorly determined, and the controlled estimate can be negative even for nonnegative cashflows. Controls do not remove general density-approximation error or shared-sample dependence. They affect decisions only: (8) and raw forward/parity diagnostics remain unadjusted.

4.3 Sparse computation and its actual complexity

For small h, many transition weights are exponentially small. The implementation screens pairs using state-dependent density bounds and then evaluates retained weights. It keeps original

IDs through sorting or batching. A fixed absolute cutoff $\varepsilon = 10^{-12}$ is used in the reported runs. Normalized weights are invariant to a common row factor, but an absolute cutoff is not; such factors must be preserved when deciding which pairs survive. A cutoff controls individual weights, not a proved bound on total omitted cashflow mass.

In one-dimensional Black–Scholes, candidate intervals can be found in sorted log spots before exponentiation. For the current Heston fused GPU implementation, all $N(N-1)$ pairs are still scanned in bounded tiles with a cheaper upper bound before expensive evaluation. It reduces storage and expensive arithmetic but remains quadratic in pair-screening work per exercise date. It does not construct or sort a flattened dense matrix. For P product columns, the fused reduction streams weighted responses, control means, control cross-products and response/control products rather than materializing a full transition matrix. Memory is controlled through tiles and partial moment buffers, including donor-block partial sums. The saved path bank and output moments also require storage according to their path, state and product counts.

5 A boundary-aware Heston implementation

The model specification is

$$\frac{dS_t}{S_t} = (r-q)\,dt + \sqrt{v_t}\,dW_t^S, \tag{16}$$

$$dv_t = \kappa(\theta - v_t)\,dt + \eta\sqrt{v_t}\,dW_t^v, \qquad d\langle W^S, W^v\rangle_t = \rho\,dt. \tag{17}$$

Here v is instantaneous variance, θ its long-run mean, κ its mean-reversion speed and η its volatility coefficient. This is the Heston diffusion [5], with $\kappa, \theta, \eta > 0$, $|\rho| < 1$ and the numerical step described next. Simulation begins at fixed $S_0 > 0, v_0 \geq 0$.

5.1 Variance endpoints and the conditional stock law

Define

$$a = \frac{2\kappa\theta}{\eta^2}, \quad c_\infty = \frac{\eta^2}{2\kappa}, \quad d = e^{-\kappa h}, \quad c = c_\infty(1-d), \quad m = dv. \tag{18}$$

An exact CIR endpoint has the Poisson–Gamma representation

$$L \mid v \sim \text{Poisson}(m/c), \qquad v' \mid L, v \sim \text{Gamma}(a + L, \text{scale } c). \tag{19}$$

This preserves nonnegative variance and its noncentral-Gamma boundary behavior without replacing the variance transition by a Gaussian.

The stock step is a finite-step approximation. Set

$$b = \frac{\rho}{\eta}\left(1 + \frac{\kappa h}{2}\right) - \frac{\rho^2 h}{4}, \qquad D = 1 - cb, \tag{20}$$

$$G_v(b) = \mathbb{E}[e^{bv'} \mid v] = D^{-a}\exp\left(\frac{mb}{D}\right), \qquad R(v,v') = \frac{(1-\rho^2)h}{2}(v+v'). \tag{21}$$

With an independent standard normal Z, simulate

$$\log\frac{S'}{S} = (r-q)h + bv' - \log G_v(b) - \frac{1}{2}R(v,v') + \sqrt{R(v,v')}Z. \tag{22}$$

The full noncentral-Gamma MGF, not a normalizer conditional on the sampled Poisson count, implies $\mathbb{E}[S' \mid S, v] = Se^{(r-q)h}$ analytically. There is no sample-dependent stock rescaling. The current step validates the finite endpoint-tilt moment domain, including $1 - 2cb > 0$.

The resulting joint density in physical coordinates is

$$\widetilde{p}_h(s', v' \mid s, v) = p_h^v(v' \mid v)\frac{\exp[-\{\log(s'/s) - \mu(v,v')\}^2/(2R(v,v'))]}{s'\sqrt{2\pi R(v,v')}}, \tag{23}$$

where $\mu = (r-q)h + bv' - \log G_v(b) - R/2$ and $R(v,v') > 0$. A zero residual variance would be singular with respect to this measure. The $1/s'$ Jacobian is essential. Equation (23) is the density of the simulated short step; it is not the exact finite-time joint Heston transition or an exact joint bridge. Exact CIR endpoints do not make the stock approximation exact.

The conditional moments used in (15) follow directly:

$$\mathbb{E}\left[\begin{pmatrix} S' \\ v' \\ S'v' \end{pmatrix} \middle| s, v\right] = \begin{pmatrix} F \\ m + ac \\ F\{ac/D + m/D^2\} \end{pmatrix}, \qquad F = se^{(r-q)h}. \tag{24}$$

In particular, the mixed moment is $FG_v'(b)/G_v(b)$. These are moments of (19)–(22), which is the law used by both simulation and weighting.

5.2 Stationary variance denominator and analytical cancellation

The stationary CIR density is

$$\pi(v') = \frac{(v')^{a-1}e^{-v'/c_\infty}}{\Gamma(a)c_\infty^a}. \tag{25}$$

The implemented local marginal approximation is $\widetilde{q}(s', v') \propto \pi(v')$: the stock factor is treated as locally constant in (S, v) coordinates. The target rows therefore use

$$\widetilde{w}_{ij} = \frac{\widetilde{p}_h(S_j', v_j' \mid S_i, v_i)}{\pi(v_j')}. \tag{26}$$

This does not initialize variance in stationarity, nor assert that stock and variance are independent. The true finite-time joint marginal from a fixed initial state generally differs from this expression. Its effect is the density distortion in (11), in addition to the transition approximation and truncation errors.

The common power of v' in the transition and stationary density can be canceled before evaluation. The CIR density admits

$$p_h^v(v' \mid v) = \frac{(v')^{a-1}}{\Gamma(a)c^a} e^{-(dv+v')/c}\, {}_0F_1\left(;a;\frac{dvv'}{c^2}\right).$$

Using $c = c_\infty(1-d)$ gives

$$\boxed{\frac{p_h^v(v' \mid v)}{\pi(v')} = \left(\frac{c_\infty}{c}\right)^a \exp\left[-\frac{d(v+v')}{c}\right] {}_0F_1\left(;a;\frac{dvv'}{c^2}\right).} \tag{27}$$

The ratio is symmetric in its two variance arguments, consistent with reversibility, and has finite limits as either endpoint tends to zero. Canceling the common power avoids separately forming a vanishing or diverging numerator and denominator near the boundary. It does not remove the remaining variance dependence.

Small arguments are evaluated with a positive ${}_0F_1$ series; interior arguments use a scaled-Bessel asymptotic expansion. These are numerical special-function evaluations of the exact CIR law. The current sparse upper bound requires $a \geq 1.5$, an implementation restriction rather than a restriction on the CIR process itself. The benchmark below has $a \simeq 1.9753$ and satisfies that restriction.

For the monthly examples, the first exercise date is after 21 or 42 simulation steps. Consequently, the configured four-step short-end window has already elapsed at every decision date: the stationary-variance and constant-stock approximation is used throughout. The retained short-end Euler marginal approximation is not exercised by these runs. A fixed four-step threshold alone is not a mathematical guarantee of local flatness or stationarity.

6 Worked examples and numerical evidence

The examples illustrate reuse of the conditional calculation across payoff structures and market dynamics. The tree and finite-difference comparisons test prices; they are not head-to-head benchmarks against least-squares Monte Carlo. In particular, the runtime comparison below compares two implementations of the same density-weighted calculation, not competing exercise methods.

6.1 A conditional ATM call: an exact identity check

In zero-rate, zero-dividend Black–Scholes, $dS_t = \sigma S_t\, dW_t$, consider a call observed at t_c conditional on $S_{t_c} = S_c$, with strike $K = S_c$ and expiry $T = t_c + \tau$. The option value at t_c is [4]

$$C(S_c, S_c, \tau) = S_c \left[2\Phi\left(\frac{\sigma\sqrt{\tau}}{2} \right) - 1 \right]. \tag{28}$$

For $S_0 = S_c = 100$, $\sigma = 0.20$ and $\tau = 0.5$, it equals 5.6371977797. It is the same as an unconditional ATM vanilla call started at 100 with remaining tenor 0.5; the earlier conditioning time does not enter the value once S_{t_c} is specified.

The density with respect to stock level is

$$p_u(y \mid s) = \frac{1}{y\sigma\sqrt{2\pi u}} \exp\left[-\frac{\{\log(y/s) + \frac{1}{2}\sigma^2 u\}^2}{2\sigma^2 u} \right]. \tag{29}$$

Thus a bank started from S_0 at zero gives the exact representation

$$C(S_c, K, \tau) = \mathbb{E}\left[\frac{p_h(S_{t_c+h} \mid S_c)}{p_{t_c+h}(S_{t_c+h} \mid S_0)} (S_T - K)^+ \right], \qquad 0 < h \le \tau. \tag{30}$$

Equation (30) provides the analytical validation target; no Monte Carlo convergence experiment for this example is reported here.

6.2 A surrenderable two-fixing Asian straddle

Consider $S_0 = K = 100$, $\sigma = 0.30$, $r = q = 0$, $t_c = 0.5$ and $T = 2$. At t_c the holder chooses an immediate payment $A = 22$, canceling the claim, or retains

$$H = \left| \frac{S_{t_c} + S_T}{2} - K \right|. \tag{31}$$

For $S_{t_c} = s$, put $k(s) = 2K - s$ and $\tau = 1.5$. The conditional value is

$$U(s) = \begin{cases} \frac{1}{2}\{C_{\mathrm{BS}}(s, k(s), \tau) + P_{\mathrm{BS}}(s, k(s), \tau)\}, & s < 2K, \\ s - K, & s \ge 2K. \end{cases} \tag{32}$$

The unconditional product remains path-dependent: the effective strike is fixed by the first observation, and surrender is decided before the terminal payoff is known. Its price is the one-dimensional benchmark

$$V_0 = \int_0^\infty \max\{A, U(s)\}\, q_{t_c}(s)\, ds. \tag{33}$$

The choice adds about 14.82% relative to retaining the claim, so the exercise feature is economically material. For target s_i, donor responses must be $H_{ij} = |(s_i + S_T^j)/2 - K|$, using the target's first fixing, with weights $p_h(S^j_{t_c+h} \mid s_i)/q_{t_c+h}(S^j_{t_c+h})$. After the decision, the booked payment is 22 or the target's own payoff $|(S^i_{t_c} + S^i_T)/2 - K|$.

Table 2: Deterministic surrender-contract benchmark. The quadrature error estimate reported by the numerical integrator is 2.22×10^{-9}.

Quantity	Value
Value with discretion	25.5731039524
Retain without discretion	22.2718231990
Always surrender	22.0000000000
Discretionary value over the better fixed choice	3.3012807534
Lower surrender boundary	80.5961870268
Upper surrender boundary	119.4038129732
Surrender probability	0.6457296344

The example also quantifies why terminal-only vanilla payoffs are insufficient. Let Z be the optimal-policy payment, carried to T at zero rate. The best square-integrable terminal-only approximation is $\mathbb{E}[Z \mid S_T]$, and its irreducible mean-square error is

$$\inf_f \mathbb{E}[(Z - f(S_T))^2] = \mathbb{E}[\mathrm{Var}(Z \mid S_T)]. \tag{34}$$

The archived conditional-moment quadrature gives RMS error 10.24662612; the residual variance is 54.05% of total cashflow variance. This bounds all terminal-only static payoff approximations, including portfolios of terminal-expiry vanillas. It says nothing comparable about dynamic hedging or portfolios spanning both fixing dates. Table 2 is a deterministic worked example, not a reported Monte Carlo validation of the current marginal-density implementation.

6.3 A dividend-paying call under Black–Scholes

For an exercise-pricing test with a nontrivial early-exercise incentive, take

$$S_0 = K = 100, \quad \sigma = 0.30, \quad r = 0, \quad q = 0.05, \quad T = 2.$$

The saved calculation uses $N = 262{,}144$ paths and 128 equal exercise intervals. GBM transitions and the marginal endpoint density are exact. The normalized nonself weights use a forward control in the exercise estimate only. A CRR tree provides both a reference on the same exercise grid and a finer approximation to continuous American exercise.

Table 3: Black–Scholes call: one Monte Carlo sample on 128 exercise intervals.

Quantity	Value
Density-weighted exercise MC	12.87419388
Matched-grid tree, 8192 tree steps	12.83218231
MC minus matched-grid tree	+0.04201157
American tree, exercise at all 8192 tree steps	12.84199017
European closed form	11.67747706
European MC on the same sample	11.71453760
Realized exercise increment over same-sample European	1.15965628
Descriptive selected-cashflow SE	0.03684673
Early exercise fraction	0.34874344

The matched-grid tree values at 2048, 4096 and 8192 tree steps are 12.83179947, 12.83203795 and 12.83218231, respectively. The MC comparison should use this finite-grid reference; the continuous-exercise American tree represents a different exercise opportunity set. The raw terminal stock forward error is +0.07943660, with raw MC standard error 0.07845415. The European payoff SE is 0.05102595. These raw simulation diagnostics are not shifted by the decision

control. A single sample and the descriptive cashflow SE do not establish statistical significance of the pricing difference or a path-count convergence rate.

6.4 Monthly Heston Bermudan puts

For (17), take

$$K = 10,\ T = 1,\ r = 0.02,\ q = 0,\ \kappa = 5,\ \theta = 0.16,\ \eta = 0.9,\ \rho = 0.1,\ v_0 = 0.15. \tag{35}$$

Exercise is permitted at $1/12, 2/12, \ldots, 1$, with no exercise at time zero. Published finite-difference values for this specification are 1.6712, 1.4507 and 1.2565 at initial spots 9.5, 10 and 10.5 [6]. They are four-decimal numerical references, not exact prices.

The saved calculation uses one bank of 65,536 paths per simulation grid. There are 21 or 42 simulation steps per monthly interval, giving 252 or 504 steps in total. The three spot cases share scaled stock paths within each grid and are not independent replications. All use (26), the controls in (24), float64 arithmetic and the fused GPU reduction.

Table 4: Heston Bermudan prices, 65,536 paths per grid. SE denotes the descriptive selected-cashflow statistic, not a confidence interval.

Simulation steps	S_0	MC price	FD reference	Difference	SE
252	9.5	1.665144	1.6712	−0.006056	0.006443
252	10.0	1.445673	1.4507	−0.005027	0.006202
252	10.5	1.252578	1.2565	−0.003922	0.005917
504	9.5	1.664578	1.6712	−0.006622	0.006386
504	10.0	1.448005	1.4507	−0.002695	0.006176
504	10.5	1.255852	1.2565	−0.000648	0.005900

At $S_0 = 10$, the corresponding European MC values are 1.43814684 and 1.44253982, leaving realized exercise contributions 0.00752662 and 0.00546522. The European reference is 1.43992607; the published Bermudan reference therefore implies an exercise premium of about 0.01077393. Comparison of total prices alone can obscure error in this smaller discretionary contribution. The raw terminal forward errors are −0.00328684 and +0.00468601, with raw standard errors 0.01681800 and 0.01693196. Appendix A gives the monthly diagnostics.

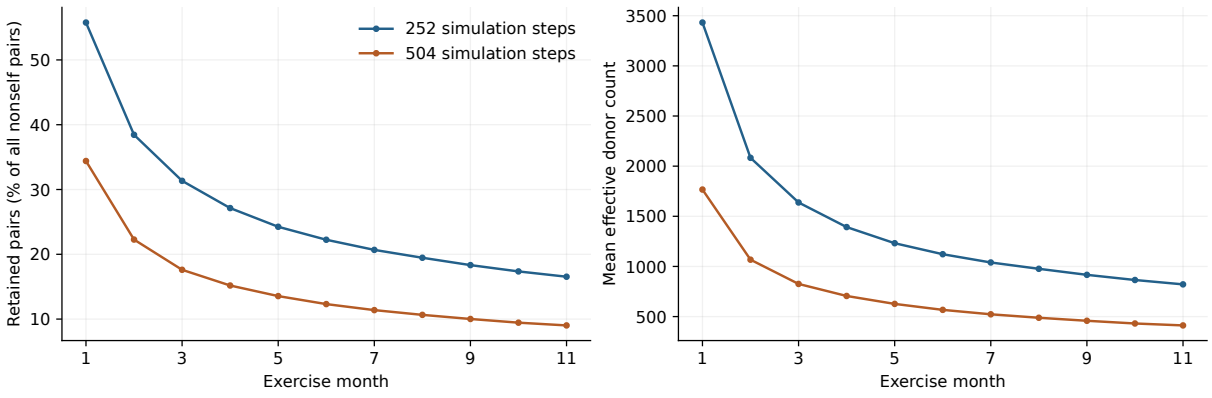


Figure 1: Saved Heston donor diagnostics at the 11 nonterminal exercise dates. Both grids have $N = 65{,}536$ paths and the same monthly exercise schedule. Effective donor count is $1/\sum_j \alpha_{ij}^2$, averaged over target rows; it is not an independent-sample size for the final estimator.

Retained pairs range from 16.5% to 55.8% of nonself pairs on the 252-step grid and from 9.0% to 34.4% on the 504-step grid. Average effective donor counts range from 821 to 3432 and from 412 to 1767, respectively. Some individual rows have effective count one. Across

the 11 dates, per contract, 6 and 18 rows require the full-donor fallback, and 15 and 76 rows have deficient three-control regression rank. These counts are identical across the three scaled contracts. Thus the finer physical grid narrows donor support as well as changing the transition approximation. It improves the observed price difference at two of the three spots but does not produce a monotone improvement at all spots. These two grids are an exploratory comparison, not a demonstrated time-step or path-count convergence study.

6.5 Computational realization and contractual integration

On an AMD Radeon RX 7900 XT, archived backward-stage timings for the three contracts together are shown in Table 5. The comparison reuses identical path banks and the same mathematical reduction. The fused timings include the recorded capture overhead and any compilation cost incurred during that execution; they are not universal hardware performance claims.

Table 5: Saved-sample backward-stage timings in seconds.

Simulation steps	Torch reductions	Fused reductions	Ratio
252	1178.414	113.747	10.36
504	823.061	94.582	8.70
Combined	2001.475	208.329	9.61

All exercise indices agree between these reductions on the 252-step sample. On the 504-step sample, one path changes exercise for each of the three scaled contracts in an ill-conditioned, three-donor regression. The largest price difference is 1.9953×10^{-5}. This is a concrete example of small changes in floating-point reduction affecting a discrete exercise decision; exact decision equivalence is not asserted for both grids.

The SHTAR integration realizes the model-independent exercise interface by separating the contractual exercise/payment processor from the numerical provider. The processor consumes decisions and books payments without specifying regressors in the provider's internal market state. It currently binds cash-settled Bermudan calls and puts with explicit exercise dates and same-day settlement. The provider supplies full-population decisions; forward processing books actual payments for the selected paths. On the same two saved banks, this integration agrees with all exercise indices of the fused implementation. Maximum individual discounted-payment differences are 5.33×10^{-15} and 4.44×10^{-15}, attributable to discounting arithmetic. This verifies payment integration, not accuracy against the optimal stopping problem.

7 Scope of the conclusions

The principal advantage over a continuation-regression formulation is generality at the product/model interface. Once a model supplies compatible conditional weights, a common engine can compare contractual alternatives without a new choice of continuation variables and basis functions for each application. Contract authors retain responsibility for payments and relevant history; model providers retain responsibility for the conditional dynamics. This separation is particularly useful when internal volatility factors or the market model change while the contract remains the same.

The exact identity establishes a way to recover a conditional expectation without waiting for a simulated path to hit the target state. It also fixes which density must appear in the denominator. The exercise implementation combines that principle with normalized estimation, approximate local marginals, finite-step transitions, moment controls and a recursively fitted policy. Each component has a separate numerical or statistical effect.

The reported examples establish an exact analytical target for conditional pricing, a materially discretionary path-dependent benchmark, a finite-grid Black–Scholes comparison, and a Heston benchmark comparison with explicit runtime and donor-support diagnostics. They support the feasibility of the implementation, rather than a claim of uniformly better prices or speed than regression-based exercise. They do not establish finite-sample unbiasedness, a rigorous optimal-price bound, a convergence rate, or a continuous-exercise limit. In particular, the monthly Heston tests retain the same 12 exercise dates when the physical simulation grid is refined. The stationary denominator does not replace the simulated variance process by a stationary process. The stock approximation and conditional moments must continue to refer to the same numerical law.

The central computational separation is maintained throughout: conditional estimates select actions, while each original path supplies the realized payments used in the price. This makes the policy estimator compatible with an explicit contractual cashflow processor while keeping the approximation and validation questions visible.

A Raw Heston forward and parity diagnostics

The following diagnostics use the unadjusted $S_0 = 10$ sample at every monthly date. The forward error is $\overline{S}_t - S_0 e^{(r-q)t}$, its SE is the ordinary stock-sample standard error, and the discounted put–call parity error is $e^{-rt}\overline{S}_t - S_0 e^{-qt}$. These quantities precede the exercise decision controls. The variance paths and stock paths are not empirically centered or rescaled.

	252 simulation steps			504 simulation steps		
Month	Forward error	Forward SE	Parity error	Forward error	Forward SE	Parity error
1	-0.001888	0.004432	-0.001884	-0.003596	0.004443	-0.003590
2	-0.004691	0.006359	-0.004676	-0.002148	0.006359	-0.002141
3	-0.004877	0.007863	-0.004852	-0.000933	0.007903	-0.000928
4	-0.002316	0.009193	-0.002301	-0.002310	0.009172	-0.002295
5	-0.001913	0.010410	-0.001897	+0.002165	0.010338	+0.002147
6	-0.001582	0.011493	-0.001566	-0.001004	0.011436	-0.000994
7	+0.001038	0.012493	+0.001026	+0.003975	0.012495	+0.003929
8	-0.004210	0.013392	-0.004155	-0.001810	0.013456	-0.001786
9	-0.006158	0.014243	-0.006066	+0.000226	0.014372	+0.000222
10	+0.002084	0.015163	+0.002050	+0.006782	0.015213	+0.006670
11	-0.000001	0.015995	-0.000001	+0.007728	0.016091	+0.007587
12	-0.003287	0.016818	-0.003222	+0.004686	0.016932	+0.004593

B Reproducibility and numerical provenance

The manuscript's companion `evidence_manifest.json` identifies the input files and SHA-256 hashes. Its build script creates the diagnostic figure and table from archived CSV files and compiles the document. It does not simulate paths or rerun pricing. Numerical sources are organized as follows:

- The conditional ATM value follows (28). The surrender benchmark and conditional terminal-payoff moments are recorded in `results_compound_full_set_gpu/report.json`, with the deterministic formulas in `compound_product.py`. Only its deterministic benchmark is used here.
- The Black–Scholes marginal-density calculation is recorded in `results_american_spot_pdf/`, including settings, prices, raw diagnostics and tree references.
- Heston prices, timing comparisons and transition diagnostics come from `results_heston_stationary_fused_gpu_65536/`. There is one Monte Carlo sample per grid; the fused reduction reuses those samples.
- Contractual integration agreement is recorded in the SHTAR project's `pricing_results/heston_discretionary_65536/validation.json`.

The reported numerical runs are research calculations at their stated path counts and control policies. No independent-repeat sampling error is estimated from the three spot-scaled contracts or from the two physical grids.

References

[1] M. Broadie and P. Glasserman (2004). A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance*, 7(4), 35–72. Author-hosted published paper.

[2] F. A. Longstaff and E. S. Schwartz (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1), 113–147. doi:10.1093/rfs/14.1.113.

[3] J. N. Tsitsiklis and B. Van Roy (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4), 694–703. doi:10.1109/72.935083. Author-hosted published paper.

[4] F. Black and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654. doi:10.1086/260062.

[5] S. L. Heston (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2), 327–343. doi:10.1093/rfs/6.2.327.

[6] D. Farahany, K. R. Jackson and S. Jaimungal (2020). Mixing LSMC and PDE methods to price Bermudan options. *SIAM Journal on Financial Mathematics*, 11(1), 201–239. doi:10.1137/19M1249035. Author preprint, arXiv:1803.07216v3; model and benchmark tables 2–3 and 9–10.