Interest Rate Models — Solutions to Problem Set 8

Section 3.3: Continuous Dynamics, Exact Fitting, and Trinomial Tree Calibration

Part A: Analytical Dynamics, Calibration, and Closed-Form Options

Problem A-1. SDE Integration and Initial Drift Extraction

Solution. (a) Define the transformed short-rate variable $f(t,r) = r(t)e^{at}$. Its partial derivatives are given by $\frac{\partial f}{\partial t} = ar(t)e^{at}$ and $\frac{\partial f}{\partial r} = e^{at}$, with a vanishing second derivative $\frac{\partial^2 f}{\partial r^2} = 0$. Applying Itô's lemma reveals:

$$df(t,r) = \frac{\partial f}{\partial t}dt + \frac{\partial f}{\partial r}dr(t) + \frac{1}{2}\frac{\partial^2 f}{\partial r^2}(dr(t))^2 = ar(t)e^{at}dt + e^{at}\left[(\vartheta(t) - ar(t))\,dt + \sigma dW(t)\right].$$

Expanding this expression shows that the terms involving $ar(t)e^{at}dt$ cancel out identically, leaving:

$$df(t,r) = e^{at}\vartheta(t)dt + \sigma e^{at}dW(t).$$

Integrating this simplified differential system explicitly over the finite horizon $[s,t]$ yields:

$$f(t,r) - f(s,r) = \int_s^t e^{au}\vartheta(u)du + \sigma\int_s^t e^{au}dW(u) \implies r(t)e^{at} = r(s)e^{as} + \int_s^t e^{au}\vartheta(u)du + \sigma\int_s^t e^{au}dW(u).$$

Multiplying the entire integrated expression by e^{-at} isolates the instantaneous short rate, verifying:

$$r(t) = r(s)e^{-a(t-s)} + \int_s^t e^{-a(t-u)}\vartheta(u)\,du + \sigma\int_s^t e^{-a(t-u)}\,dW(u).$$

(b) The stochastic integral term $I_t = \sigma\int_s^t e^{-a(t-u)}dW(u)$ constitutes a continuous martingale under the filtration $\mathcal{F}_s$. Consequently, its conditional expectation vanishes identically: $\mathbb{E}[I_t \mid \mathcal{F}_s] = 0$. Taking the conditional expectation of the isolated short rate reveals that the remaining components are entirely deterministic under $\mathcal{F}_s$:

$$\mathbb{E}[r(t) \mid \mathcal{F}_s] = r(s)e^{-a(t-s)} + \int_s^t e^{-a(t-u)}\vartheta(u)\,du.$$

To evaluate the conditional variance, note that because $r(s)$ and the drift integral are deterministic at time s, they do not contribute to the variance. Applying the Itô isometry to the remaining stochastic integral gives:

$$\mathrm{Var}(r(t) \mid \mathcal{F}_s) = \mathbb{E}\left[\left(\sigma\int_s^t e^{-a(t-u)}dW(u)\right)^2 \,\middle|\, \mathcal{F}_s\right] = \sigma^2\int_s^t e^{-2a(t-u)}du.$$

Using the variable transformation $v = t-u$, where $dv = -du$, the boundary integration limits reverse from $[s,t]$ to $[t-s, 0]$:

$$\mathrm{Var}(r(t) \mid \mathcal{F}_s) = \sigma^2\int_0^{t-s} e^{-2av}dv = \left[-\frac{\sigma^2}{2a}e^{-2av}\right]_0^{t-s} = \frac{\sigma^2}{2a}\left(1 - e^{-2a(t-s)}\right).$$

(c) Setting the valuation starting point to $s = 0$ and substituting the identity $r(t) = x(t) + \alpha(t)$, where $\mathbb{E}[x(t)] = 0$, implies that $\mathbb{E}[r(t)] = \alpha(t)$. Differentiating the expected value of our integrated short rate with respect to maturity t gives:

$$\frac{\partial \alpha(t)}{\partial t} = \vartheta(t) - a \int_0^t e^{-a(t-u)} \vartheta(u) du = \vartheta(t) - a \left[\alpha(t) - r(0)e^{-at}\right].$$

Using the analytical expression for the market instantaneous forward rate, $f^M(0,t) = \alpha(t) - \frac{\sigma^2}{2a^2}(1 - e^{-at})^2$, we take its partial derivative with respect to time to obtain:

$$\frac{\partial f^M(0,t)}{\partial t} = \frac{\partial \alpha(t)}{\partial t} - \frac{\sigma^2}{a}(1 - e^{-at})e^{-at}.$$

Isolating $\vartheta(t)$ by substituting these relations back into the localized differentiated system yields:

$$\vartheta(t) = \frac{\partial f^M(0,t)}{\partial t} + af^M(0,t) + \frac{\sigma^2}{2a}\left(1 - e^{-2at}\right).$$

(d) Because the continuous process $r(t)$ is normally distributed, its absolute density spans the entire real line ($r \in (-\infty, \infty)$). The risk-neutral probability of generating a negative short rate is found by standardizing the boundary variable:

$$\mathbb{Q}\{r(t) < 0\} = \mathbb{Q}\left(\frac{r(t) - \alpha(t)}{\sqrt{\text{Var}(r(t))}} < \frac{-\alpha(t)}{\sqrt{\text{Var}(r(t))}}\right) = \Phi\left(\frac{-\alpha(t)}{\sqrt{\frac{\sigma^2}{2a}(1 - e^{-2at})}}\right).$$

This structural feature is the main limitation of the model because negative interest rates allow for arbitrage in the presence of physical cash hoarding (which carries a floor of 0%). This can lead to distortion when pricing long-dated out-of-the-money options.

Problem A-2. The T-Forward Measure and Zero-Coupon Bond Options

Solution. (a) The Radon–Nikodym derivative that governs the change of numeraire from the risk-neutral bank account $B(t) = \exp(\int_0^t r(u)du)$ to the T-forward discount bond numeraire asset $P(t,T)$ is given by the ratio of their relative values scaled to today:

$$\frac{d\mathbb{Q}^T}{d\mathbb{Q}} = \frac{P(T,T)/P(0,T)}{B(T)/B(0)} = \frac{1}{P(0,T)\exp\left(\int_0^T r(u)du\right)} = \frac{\exp\left(-\int_0^T r(u)du\right)}{P(0,T)}.$$

(b) Under the risk-neutral measure $\mathbb{Q}$, the zero-coupon bond price has an absolute volatility vector given by its diffusion coefficient. By applying Itô's lemma to the affine pricing equation $P(t,T) = A(t,T)e^{-B(t,T)r(t)}$, its diffusion coefficient is isolated as $-B(t,T)\sigma P(t,T)$. The Girsanov theorem dictates that changing the numeraire asset from $B(t)$ to $P(t,T)$ alters the underlying drift of any stochastic process by subtracting the volatility of the new numeraire. For the process $x(t)$, this transformation gives:

$$dx(t) = -ax(t)dt + \sigma\left(dW^T(t) - B(t,T)\sigma dt\right) = \left[-B(t,T)\sigma^2 - ax(t)\right]dt + \sigma dW^T(t).$$

(c) Because the numeraire asset matches the expiration date of the contract ($P(T,T) = 1$), changing to the T-forward measure pulls the stochastic discount factor completely outside the expectation operator, simplifying the expression to:

$$\mathrm{ZBC}(t,T,S,X) = P(t,T)\,\mathbb{E}^T\left[(P(T,S) - X)^+ \mid \mathcal{F}_t\right].$$

Under this forward measure, the log-bond price $\ln P(T,S)$ is normally distributed. Evaluating this expectation is equivalent to computing a standard Black option price where the underlying asset is the forward bond price $\frac{P(t,S)}{P(t,T)}$. Integrating the total variance accumulated by the option over the horizon $[t,T]$ isolates the integrated volatility parameter:

$$\sigma_p = \sqrt{\int_t^T \sigma^2 e^{-2a(T-u)} B(T,S)^2 du} = \sigma\sqrt{\frac{1-e^{-2a(T-t)}}{2a}}B(T,S).$$

This parameter acts as the standardized scaling factor for the asset's variance. Substituting this variance into the lognormal distribution results in the closed-form Black-type equation:

$$\mathrm{ZBC}(t,T,S,X) = P(t,S)\Phi(h) - XP(t,T)\Phi(h-\sigma_p), \quad \text{where } h = \frac{1}{\sigma_p}\log\left(\frac{P(t,S)}{P(t,T)X}\right) + \frac{\sigma_p}{2}.$$

Here, h functions as a standardized Z-score that normalizes the financial distance to the strike hurdle relative to the model's volatility.

(d) Put-call parity for options written on zero-coupon bonds requires that $\mathrm{ZBC}(t,T,S,X) - \mathrm{ZBP}(t,T,S,X) = P(t,S) - XP(t,T)$. Substituting the explicit pricing equation for ZBC and utilizing the identity $1 - \Phi(h) = \Phi(-h)$ isolates the closed-form formula for the European put:

$$\mathrm{ZBP}(t,T,S,X) = XP(t,T)\Phi(-h+\sigma_p) - P(t,S)\Phi(-h).$$

Problem A-3. Jamshidian's Decomposition for Coupon Options

Solution. (a) Pricing a portfolio-level option is mathematically difficult because the underlying asset is a linear combination of separate lognormal random variables ($\sum c_i P(T, t_i)$). The sum of individual lognormal distributions does not follow a known, simple distribution. Consequently, the multi-cashflow option cannot be solved with a simple closed-form integration and typically requires complex multi-dimensional simulations.

(b) In a one-factor short-rate model, all asset uncertainty across the entire timeline is driven by a single source of randomness ($r(t)$). Under the affine pricing structure, each individual component zero-coupon bond is defined as $P(T, t_i) = A(T, t_i)e^{-B(T,t_i)r(T)}$. Taking the partial derivative with respect to the state variable yields:

$$\frac{\partial P(T, t_i)}{\partial r(T)} = -B(T, t_i)P(T, t_i) < 0 \quad (\text{since } B(T, t_i) > 0 \text{ for } t_i > T).$$

Because every single bond option component has a strictly negative first derivative with respect to $r(T)$, any linear portfolio combination of these components is a strictly decreasing, monotonic function of the short rate.

(c) Because the underlying portfolio value decreases monotonically as $r(T)$ rises, there is exactly one unique interest rate value r^* where the total market value of the coupon bond matches the contract strike hurdle X. A one-dimensional root-finding algorithm (such as Newton-Raphson) can isolate this critical rate by solving:

$$f(r^*) = \sum_{i=1}^{n} c_i A(T, t_i)e^{-B(T,t_i)r^*} - X = 0.$$

The algorithm updates iteratively via $r^*_{new} = r^*_{old} - \frac{f(r^*_{old})}{f'(r^*_{old})}$, where $f'(r^*) = -\sum c_i B(T, t_i)A(T, t_i)e^{-B(T,t_i)r^*}$, until it converges to the exact root.

(d) At the option's expiration date T, if the realized short rate satisfies $r(T) > r^*$, the monotonicity property guarantees that the total portfolio value will fall below the strike price X, putting the contract in-the-money:

$$\left(X - \sum_{i=1}^{n} c_i P(T, t_i)\right)^+ = \sum_{i=1}^{n} c_i \left(A(T, t_i)e^{-B(T,t_i)r^*} - P(T, t_i)\right)^+ = \sum_{i=1}^{n} c_i (X_i - P(T, t_i))^+.$$

Because the exercise boundary for the entire portfolio matches the individual exercise boundary for every coupon payment, the consolidated payoff splits into separate independent payoffs. Taking expectations confirms that the swaption or coupon bond option can be written as a simple sum of individual zero-coupon bond puts:

$$\text{PS}(t, T, \mathcal{T}, N, X) = N \sum_{i=1}^{n} c_i \, \text{ZBP}(t, T, t_i, X_i), \qquad \text{where } X_i = A(T, t_i)e^{-B(T,t_i)r^*}.$$

Part B: Discrete Numerical Lattices and Arrow–Debreu Calibration

Problem B-1. Mean Reversion and Dynamic Branch Selection

Solution. (a) Integrating the zero-mean Ornstein–Uhlenbeck SDE over a single horizontal step Δt_i establishes the continuous local targets: $M_{i,j} = x_{i,j}e^{-a\Delta t_i}$ and $V_i^2 = \frac{\sigma^2}{2a}(1 - e^{-2a\Delta t_i})$. The vertical space interval is fixed to the scale $\Delta x_i = V_{i-1}\sqrt{3}$ based on standard numerical analysis principles. Setting the grid spacing proportional to the standard deviation scaled by $\sqrt{3}$ maximizes the local accuracy of the Taylor series expansion, minimizes truncation errors, and ensures that the calculated branch probabilities remain stable and positive.

(b) Under a naive lattice structure that only allows horizontal branching paths, the middle branch destination from node (i, j) is locked into the same vertical index j at the next time step ($x_{i+1,center} = j\Delta x_{i+1}$). If a random path drifts far above the long-term mean, the continuous mean target $M_{i,j}$ will be pulled downward by mean reversion. Eventually, at high or low states, $M_{i,j}$ will drift entirely past the adjacent rows ($j + 1$ or $j - 1$). If the physical rails do not shift to follow this pull, the algebraic equations will generate probabilities greater than 100% or less than 0% to force a match, causing the numerical system to break down.

(c) The dynamic branching indicator variable $k = \text{round}\left(\frac{M_{i,j}}{\Delta x_{i+1}}\right)$ calculates the exact theoretical mean target, divides it by the vertical step size, and rounds the result to the nearest integer index. This integer identifies the physical row coordinate that sits closest to where mean reversion wants to pull the process. By centering the three branching options $(k + 1, k, k - 1)$ around this updated index rather than the original position j, the tree actively tilts its branches to follow the mean, keeping the true target inside its variance window and ensuring the probabilities remain stable and well-behaved.

(d) Three distinct branching topologies can emerge across the grid layout:

1. **Standard Symmetric Branching (Type B):** Occurs near the center of the tree where mean reversion is weak. The rounding factor matches the current row ($k = j$), and the paths branch standardly to $j + 1$, j, and $j - 1$.
2. **Down-Shifting Branching (Type C):** Occurs at high interest rate states. Strong mean reversion pulls the expected path down, so $k < j$. The three branching lines bend downward to target $k + 1$, k, and $k - 1$.
3. **Up-Shifting Branching (Type A):** Occurs at low interest rate states. The process is pulled upward, so $k > j$. The three branching lines tilt upward to target $k + 1$, k, and $k - 1$.

Problem B-2. Transition Probability Constraints

Solution. (a) To calibrate the local transition weights (p_u, p_m, p_d), the discrete steps must match the total probability condition, preserve the local mean target, and align with the local variance target:

$$\begin{cases} p_u + p_m + p_d = 1 \\ p_u(x_{i+1,k+1}) + p_m(x_{i+1,k}) + p_d(x_{i+1,k-1}) = M_{i,j} \\ p_u(x_{i+1,k+1} - M_{i,j})^2 + p_m(x_{i+1,k} - M_{i,j})^2 + p_d(x_{i+1,k-1} - M_{i,j})^2 = V_i^2 \end{cases}$$

Substituting the grid definition $x_{i+1,k} = k\Delta x_{i+1}$ and the spatial error term $\eta_{j,k} = M_{i,j} - k\Delta x_{i+1}$ simplifies the system to:

$$\begin{cases} p_u + p_m + p_d = 1 \\ (p_u - p_d)\Delta x_{i+1} = \eta_{j,k} \\ (p_u + p_d)\Delta x_{i+1}^2 = V_i^2 + \eta_{j,k}^2 \end{cases}$$

(b) We isolate p_u and p_d by combining the second and third conditions of our simplified system:

$$p_u + p_d = \frac{V_i^2 + \eta_{j,k}^2}{\Delta x_{i+1}^2}, \qquad p_u - p_d = \frac{\eta_{j,k}}{\Delta x_{i+1}}.$$

Adding and subtracting these equations isolates the individual probabilities. Then, substituting the vertical grid spacing parameter $\Delta x_{i+1} = V_i\sqrt{3}$ (which means $\Delta x_{i+1}^2 = 3V_i^2$) gives:

$$p_u = \frac{1}{6} + \frac{\eta_{j,k}^2}{6V_i^2} + \frac{\eta_{j,k}}{2\sqrt{3}V_i}, \qquad p_d = \frac{1}{6} + \frac{\eta_{j,k}^2}{6V_i^2} - \frac{\eta_{j,k}}{2\sqrt{3}V_i}.$$

Finally, we solve for the middle probability p_m using the total probability condition ($p_m = 1 - p_u - p_d$):

$$p_m = 1 - \left(\frac{2}{6} + \frac{2\eta_{j,k}^2}{6V_i^2}\right) = \frac{2}{3} - \frac{\eta_{j,k}^2}{3V_i^2}.$$

(c) For p_m to remain a valid probability, it cannot drop below zero ($p_m \geq 0$). Setting the algebraic solution to this lower bound gives:

$$\frac{2}{3} - \frac{\eta_{j,k}^2}{3V_i^2} \geq 0 \implies \frac{\eta_{j,k}^2}{3V_i^2} \leq \frac{2}{3} \implies \eta_{j,k}^2 \leq 2V_i^2 \implies |\eta_{j,k}| \leq V_i\sqrt{2}.$$

Crucially, checking the boundaries for the outer branching tracks ($p_u \geq 0$ and $p_d \geq 0$) imposes an even tighter constraint:

$$\frac{1}{6} + \frac{\eta_{j,k}^2}{6V_i^2} - \frac{|\eta_{j,k}|}{2\sqrt{3}V_i} \geq 0.$$

Solving this quadratic inequality reveals that for all three probabilities to remain strictly positive and bounded between 0 and 1, the maximum spatial alignment error cannot exceed:

$$|\eta_{j,k}| \leq \frac{V_i}{\sqrt{3}}.$$

(d) Because the vertical spacing between adjacent grid rails is exactly $\Delta x_{i+1} = V_i\sqrt{3}$, the maximum possible distance between any continuous point and the nearest discrete row rung is exactly half of that grid step: $\frac{1}{2}\Delta x_{i+1} = \frac{V_i\sqrt{3}}{2}$. Using the integer rounding function $k = \text{round}(M_{i,j}/\Delta x_{i+1})$ guarantees that the spatial error $\eta_{j,k}$ will never exceed this mid-point threshold:

$$|\eta_{j,k}| \leq \frac{1}{2}\Delta x_{i+1} = \frac{\sqrt{3}}{2}V_i.$$

Comparing this guaranteed maximum error against the probability stability limit ($|\eta_{j,k}| \leq \frac{1}{\sqrt{3}}V_i \approx 0.577V_i$) confirms that the rounding mechanism keeps the system well within its safe operating bounds, validating the structural integrity of the lattice.

Problem B-3. Arrow–Debreu Forward Induction Loops

Solution. (a) The Arrow–Debreu price $Q_{i+1,j}$ for a node in the next column is calculated by summing the values of all paths leading into it from the current column. Each incoming step multiplies the prior node value ($Q_{i,h}$) by its specific transition probability ($q(h, j)$) and discounts the cash flow using the real interest rate mapped to that originating node ($r_{i,h} = \alpha_i + h\Delta x_i$):

$$Q_{i+1,j} = \sum_h Q_{i,h} q(h, j) \exp\left(-\left[\alpha_i + h\Delta x_i\right] \Delta t_i\right).$$

(b) Derivative options are priced using a backward recursive loop because their payoffs are tied to a known boundary condition at expiration (T). The engine sets those terminal values first and works backward to find today's price. In contrast, Arrow–Debreu prices are calibrated using a forward induction loop because the initial state today is known with certainty ($Q_{0,0} = 1.0$). The engine works forward across time to build the probability and discounting matrix step-by-step. This forward pass must occur first to define the grid before any option contracts can be evaluated.

(c) A zero-coupon market bond that promises to pay exactly \$1 at time step t_{i+1} can be replicated on our tree by purchasing the entire portfolio of Arrow–Debreu lottery tickets for that column. Since the train must land on one of the nodes, the sum of these discounted states must match the actual market bond price $P^M(0, t_{i+1})$:

$$P^M(0, t_{i+1}) = \sum_{j=\underline{j}_i}^{\bar{j}_i} Q_{i,j} \exp\left(-r_{i,j}\Delta t_i\right) = \sum_{j=\underline{j}_i}^{\bar{j}_i} Q_{i,j} \exp\left(-[\alpha_i + j\Delta x_i]\Delta t_i\right).$$

Cleanly factoring the constant column vector element $\exp(-\alpha_i\Delta t_i)$ out of the summation layer reveals:

$$P^M(0, t_{i+1}) = \exp(-\alpha_i\Delta t_i) \sum_{j=\underline{j}_i}^{\bar{j}_i} Q_{i,j} \exp\left(-j\Delta x_i\Delta t_i\right).$$

Taking the natural logarithm of both sides isolates the calibration parameter, verifying the master shift equation:

$$\ln P^M(0, t_{i+1}) = -\alpha_i\Delta t_i + \ln\left(\sum_{j=\underline{j}_i}^{\bar{j}_i} Q_{i,j} \exp\left(-j\Delta x_i\Delta t_i\right)\right) \implies \alpha_i = \frac{1}{\Delta t_i} \log\left(\frac{\sum_{j=\underline{j}_i}^{\bar{j}_i} Q_{i,j} \exp(-j\Delta x_i\Delta t_i)}{P^M(0, t_{i+1})}\right).$$

(d) This two-stage staging approach provides a significant computational advantage. Because the zero-mean process $x(t)$ depends only on the constant parameters a and σ, the core tree geometry, row connections (k), and branch probabilities (p_u, p_m, p_d) are entirely stable. They only need to be computed once. If market bond prices shift tomorrow morning, the underlying structure of the tree remains completely unchanged. The software engine simply reruns the forward loop to update the vertical shifts (α_i). This separates the stable probability model from erratic market data, maximizing performance and ensuring the lattice remains robust.