

MONOTONE 2D INTEGRATION SCHEME FOR MEAN-CVaR OPTIMIZATION VIA FOURIER-TRAINED TRANSITION KERNELS

DUY-MINH DANG* AND HAO ZHOU †

Abstract. We present a strictly monotone, provably convergent two-dimensional (2D) integration method for multi-period mean-conditional value-at-risk (mean-CVaR) reward-risk stochastic control in models whose one-step increment law is specified via a closed-form characteristic function (CF). When the transition density is unavailable in closed form, we learn a nonnegative, normalized 2D transition kernel in Fourier space using a simplex-constrained Gaussian-mixture parameterization, and discretize the resulting convolution integrals with composite quadrature rules with nonnegative weights to guarantee monotonicity. The scheme is implemented efficiently using 2D fast Fourier transforms. Under mild Fourier-tail decay assumptions on the CF, we derive Fourier-domain L_2 kernel-approximation and truncation error estimates and translate them into real-space bounds that are used to establish ℓ_∞ -stability, consistency, and pointwise convergence as the discretization and kernel-approximation parameters vanish. Numerical experiments for a fully coupled 2D jump-diffusion model in a multi-period portfolio optimization setting illustrate robustness and accuracy.

Key words. reward-risk stochastic control, mean-CVaR, monotone schemes, characteristic functions, Fourier-trained transition kernels

MSC codes. 65M12, 65T50, 93E20, 60E10

1 Introduction Reward-risk optimization arises in many settings where decisions trade off performance against adverse outcomes. Such formulations appear in environmental and resource management [19], supply-chain planning [17], and more broadly in engineering and applied decision-making [13]. In finance, the same trade-off is central in multi-period portfolio choice and retirement saving under market and macroeconomic uncertainty [14].

A widely used tail-risk measure in such formulations is conditional value-at-risk (CVaR), also known as expected shortfall [28]. It quantifies tail risk via a conditional expectation beyond a quantile threshold and admits a threshold-based representation that is computationally convenient. In reward-risk optimization, the reward is typically an expected gain or payoff, making it natural to pair this mean term with a CVaR penalty. The resulting mean-CVaR objective is widely used in multi-period stochastic control [14, 25].

From a numerical perspective, value functions arising in stochastic control, including those induced by mean-CVaR criteria, are typically nonsmooth. Since optimal decisions are obtained by comparing numerical value functions across admissible controls, convergence of the discretization is essential. In particular, if the scheme is not monotone, the computed value function may converge to the wrong solution, leading to unreliable policies [2, 26]. This motivates the design of monotone schemes for reward-risk control problems.

Monotone schemes are central to the convergence theory of Barles-Souganidis [2]; however, their construction and analysis are most fully developed in one-dimensional (1D) models, both with and without jumps. Constructing strictly monotone schemes in multiple dimensions is substantially more challenging, especially under correlation and nonlocal jump effects, and robust treatments of multi-dimensional integro-

*School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, (duyminh.dang@uq.edu.au).

†School of Mathematics and Physics, The University of Queensland, St Lucia, Brisbane 4072, Australia, (h.zhou3@uq.net.au).

differential terms remain difficult [21].

In many popular stochastic models, the transition density is not available in closed form, while its Fourier transform, i.e. the characteristic function (CF) of the increments is available analytically [12]. When CFs are known, Fourier-based methods can then handle complex dynamics efficiently and provide an attractive alternative to finite differences. Among these, the Fourier-cosine (COS) method is particularly notable for achieving high-order convergence in piecewise smooth problems [12, 29]. However, for general stochastic control problems, which often involve nonsmooth structures, such convergence is typically unattainable [20]. More importantly, standard Fourier discretizations may lose monotonicity, leading to spurious value functions and unreliable policies [11].

These challenges have motivated strictly monotone integration methods that enforce nonnegative conditional density representations [33]. Like Fourier-based approaches, they leverage transform-domain information. However, existing constructions typically rely on additional structure, for example, effectively 1D dynamics or simplified multi-dimensional models with closed-form nonnegative densities, and are available only in limited 2D settings [34, 10].

This paper develops a strictly monotone and provably convergent 2D integration framework for multi-period mean-CVaR optimization in fully coupled 2D models. Our main contribution is to enforce strict monotonicity without requiring an explicit closed-form transition density, by learning a nonnegative 2D transition kernel directly from the increment law via its CF. Building on the Fourier-domain density-learning idea in [11], we approximate the kernel using a Gaussian-mixture representation fitted to the CF. Constraining the mixture weights to the probability simplex ensures nonnegativity and normalization by construction. The resulting 2D convolution integrals are discretized using composite quadrature rules with nonnegative weights, yielding a strictly monotone scheme that can be implemented efficiently using 2D FFTs.

A second contribution is a Fourier-to-real-space analysis that underpins convergence of the scheme. Leveraging the explicit Fourier form of both the Gaussian-mixture kernel and the CF, we derive Fourier-domain L_2 error estimates and translate them into real-space bounds. Under mild Fourier-tail decay assumptions, these bounds are used to establish ℓ_∞ -stability, consistency, and pointwise convergence as the discretization and kernel-approximation parameters vanish.

We keep the mean-CVaR formulation and modelling assumptions broad: between decision times the 2D increment law is specified via a closed-form CF with mild Fourier-tail decay. This setting covers a wide class of coupled 2D models where transition densities are unavailable but CFs are tractable. To illustrate the method, we consider a multi-period portfolio optimization problem calibrated to long-horizon market data. More generally, the proposed framework extends to other reward-risk stochastic control problems with discrete interventions and translation-invariant kernels.

2 Modelling We work on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{0 \leq t \leq T}, \mathbb{P})$ over a finite horizon $T > 0$.

Intervention times. We fix equally spaced decision times in $[0, T]$:

$$(2.1) \quad \mathcal{T} = \{t_m \mid t_m = m\Delta t, m = 0, \dots, M\}, \quad \Delta t = T/M,$$

with $t_0 = 0$. For any process $\{Z_t\}_{t \in [0, T]}$ with left/right limits at $\{t_m\}$, we write $Z_m^- := \lim_{\varepsilon \rightarrow 0^+} Z(t_m - \varepsilon)$ and $Z_m^+ := \lim_{\varepsilon \rightarrow 0^+} Z(t_m + \varepsilon)$.

State variables and log representation. Let A_t^s and A_t^b denote two nonnegative state components at time $t \in [0, T]$ (e.g. two resource/inventory levels, two account values, or two components of a controlled system).

Intervention actions occur only at decision times. At each t_m , $m = 0, \dots, M - 1$, the state may be updated first by a deterministic exogenous input processed at t_m^- (assumed nonnegative; e.g. an external injection/replenishment) and then by an intervention applied at t_m^+ . At $t_M = T$, no further intervention is applied and a terminal reward functional is evaluated. Between decision times, i.e. on each interval $[t_m^+, t_{m+1}^-]$, no interventions are applied and the dynamics evolve according to the (uncontrolled) law specified below.

To work in log coordinates, we impose a strictly positive floor at decision times: after the contribution/intervention at t_m , we enforce $A_{t_m^+}^s, A_{t_m^+}^b \geq e^{w_\infty}$ for fixed $w_\infty \ll 0$. Hence, for $t \in [t_m^+, t_{m+1}^-]$,

$$S_t = \ln A_t^s, \quad B_t = \ln A_t^b, \quad X_t = (S_t, B_t) \in \mathbb{R}^2.$$

We write $x = (s, b)$ for a generic state, $y = (y_s, y_b)$ for an integration variable, and $\eta = (\eta_s, \eta_b)$ for its Fourier (frequency) counterpart.

Increment law and Fourier transform. On each inter-decision interval $[t_m^+, t_{m+1}^-]$, we define the one-step increment

$$(2.2) \quad (\Delta X)_m \equiv ((\Delta S)_m, (\Delta B)_m) := X_{m+1}^- - X_m^+.$$

We assume the conditional law of $(\Delta X)_m$ given X_m^+ depends only on Δt and is translation-invariant; in particular, the law of $(\Delta X)_m$ does not depend on m . Let ΔX denote a generic one-step increment with this common law and define its CF

$$G(\eta; \Delta t) := \mathbb{E}[e^{i\eta \cdot \Delta X}], \quad \eta \in \mathbb{R}^2.$$

When ΔX admits a density $g(\cdot; \Delta t)$ on \mathbb{R}^2 , g and G form a Fourier pair:

$$(2.3) \quad G(\eta; \Delta t) = \int_{\mathbb{R}^2} e^{i\eta \cdot y} g(y; \Delta t) dy, \quad g(y; \Delta t) = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} e^{-i\eta \cdot y} G(\eta; \Delta t) d\eta.$$

ASSUMPTION 2.1 (Modelling assumptions). *For each $\Delta t > 0$, the uncontrolled one-step increment ΔX admits a density $g(\cdot; \Delta t)$ on \mathbb{R}^2 , and:*

- (A1) (*Homogeneous one-step kernel.*) *For all m and all $x \in \mathbb{R}^2$, the conditional density of X_{m+1}^- given $X_m^+ = x$ is $y \mapsto g(y - x; \Delta t)$ (time-homogeneous and translation-invariant).*
- (A2) (*Regularity and boundedness.*) $g(\cdot; \Delta t) \in L_1(\mathbb{R}^2) \cap L_\infty(\mathbb{R}^2)$.
- (A3) (*Closed-form CF.*) $G(\eta; \Delta t)$ is available in closed form for all $\eta \in \mathbb{R}^2$.
- (A4) (*Fourier tail control.*) *There exist $\alpha \in (0, 2]$, $c_0 > 0$, and $R_{\text{tail}} \geq 0$ such that $|G(\eta; \Delta t)| \leq \exp(-c_0 \Delta t \|\eta\|_2^\alpha)$ for all $\|\eta\|_2 \geq R_{\text{tail}}$.*
- (A5) (*Exponential moments.*) $G(-i, 0; \Delta t)$, $G(0, -i; \Delta t)$, and $G(-i, -i; \Delta t)$ are finite, and $G(0, 0; \Delta t) = 1$.

Remark 2.2 (Examples and scope). Assumption 2.1 covers a broad class of translation invariant increment laws with tractable CFs. In particular, it holds for many 2D exponential-Lévy and jump-diffusion models with closed-form characteristic functions, including correlated Brownian increments, Merton [24] and Kou [18] jump-diffusions with a nondegenerate Gaussian component, Normal-Inverse-Gaussian [3], tempered-stable and CGMY models with $Y \in (0, 2)$ [5], and finite Gaussian mixtures [23], including variants with co-jumps. Condition (A4) provides the decay needed for Fourier truncation, while (A5) ensures the exponential tilts used in boundary propagation are admissible.

3 Mean–CVaR optimization This section specifies a two-component reward–risk stochastic control problem with discrete interventions. The notation follows an allocation setting in which an aggregate level is redistributed between two nonnegative components. Allocating inventory between two storage locations (or water between two reservoirs) is one concrete example.

3.1 Intervention map Suppose the system is in state $x = (s, b)$ at time t_m^- (immediately before the deterministic exogenous input $q_m \geq 0$), with aggregate level $W_m^- = e^s + e^b$. Immediately after the input, at time t_m^+ , the aggregate level is

$$(3.1) \quad W_m^+ = W_m^- + q_m = e^s + e^b + q_m, \quad m = 0, \dots, M-1.$$

An admissible Markov feedback control at t_m is any Borel–measurable map

$$(3.2) \quad u_m : \mathbb{R}^2 \rightarrow \mathcal{Z}, \quad (s, b) \mapsto u_m(s, b).$$

Here, $\mathcal{Z} := [0, 1]$ is the set of admissible actions: u_m is interpreted as the fraction of W_m^+ assigned to the first component, and $1 - u_m$ to the second (a simplex-type constraint). For brevity, when the arguments are clear we write $u_m := u_m(s, b)$.

After applying u_m , the post-intervention levels of the two components are $u_m W_m^+$ and $(1 - u_m)W_m^+$, respectively, with W_m^+ given by (3.1). To work in log coordinates, introduce a strictly positive floor e^{w_∞} with fixed $w_\infty \ll 0$. The post-intervention log-state is $X_m^+ = (S_m^+, B_m^+)$, where

$$(3.3) \quad \begin{aligned} S_m^+ &:= s^+(s, b, q_m, u_m) = \ln(\max\{u_m W_m^+, e^{w_\infty}\}), \\ B_m^+ &:= b^+(s, b, q_m, u_m) = \ln(\max\{(1 - u_m)W_m^+, e^{w_\infty}\}). \end{aligned}$$

An admissible policy is $\mathcal{U}_0 = \{u_m\}_{m=0}^{M-1}$ with each u_m as in (3.2). We denote the set of such policies by \mathcal{A} , and write $\mathcal{U}_m = \{u_{m'}\}_{m'=m}^{M-1}$ for the tail policy applied from t_m onwards.

3.2 Mean–CVaR formulation Although CVaR is often defined for losses, we apply it here to the terminal outcome W_T , adopting the convention that larger values indicate better performance. Equivalently, one may apply the standard loss-based CVaR to $-W_T$. Fix $\alpha \in (0, 1)$. Following [27, 25], for any policy $\mathcal{U}_0 \in \mathcal{A}$, the CVaR of W_T at level α admits the threshold representation

$$(3.4) \quad \text{CVaR}_{\alpha, \mathcal{U}_0}^{x_0, t_0^-} = \sup_{w \geq 0} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} \left[w + \frac{1}{\alpha} \min(W_T - w, 0) \right],$$

where the feasible set $w \in [0, \infty)$ matches the attainable range of W_T .

Using a scalarization parameter $\gamma > 0$, we adopt a *pre-commitment* mean–CVaR objective: at inception t_0^- the decision maker selects a policy $\mathcal{U}_0 \in \mathcal{A}$ and a CVaR threshold $w \geq 0$ to maximize the criterion, and then commits to \mathcal{U}_0 thereafter. The pre-commitment mean–CVaR value function is

$$(3.5) \quad \mathcal{V}(x_0, t_0^-) = \sup_{\mathcal{U}_0 \in \mathcal{A}} \sup_{w \geq 0} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} \left[W_T + \gamma \left(w + \frac{1}{\alpha} \min(W_T - w, 0) \right) \right].$$

Since \mathcal{U}_0 and w range over independent sets and w enters only through the terminal objective functional, the iterated suprema commute:

$$(3.6) \quad \mathcal{V}(x_0, t_0^-) = \sup_{w \geq 0} \sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} \left[W_T + \gamma \left(w + \frac{1}{\alpha} \min(W_T - w, 0) \right) \right],$$

subject to Assumption 2.1 on modelling, the exogenous inputs (3.1), and the intervention map (3.3).

Remark 3.1 (Finiteness of $\mathbb{E}_{\mathcal{U}_0^{x_0, t_0^-}} [W_T]$). Under Assumption 2.1 (A5), we have $\mathbb{E}[e^{\Delta S}] = G(-i, 0; \Delta t) < \infty$, and $\mathbb{E}[e^{\Delta B}] = G(0, -i; \Delta t) < \infty$. Since interventions are restricted to $u \in [0, 1]$ and the exogenous inputs $\{q_m\}$ are deterministic and finite, these moment bounds propagate over the finite horizon to yield $\mathbb{E}_{\mathcal{U}_0^{x_0, t_0^-}} [W_T] < \infty$ uniformly over admissible policies $\mathcal{U}_0 \in \mathcal{A}$.

LEMMA 3.2 (Existence of a finite optimal threshold). *Fix $\gamma > 0$. Then the outer optimization in (3.6) is finite and is attained by some $w^*(x_0) \in [0, \infty)$.*

A proof of Lemma 3.2 is given in Appendix A. In particular, Lemma 3.2 (together with Remark 3.1) implies that the objective functional in (3.5)–(3.6) is finite (and hence well-defined).

Lifted-state recursion (fixed w). We reformulate (3.6) as an equivalent lifted-state problem [25]. For fixed $w \geq 0$, define the auxiliary value on the augmented state (x, w, t) , $x = (s, b)$:

$$(3.7) \quad V(x, w, t_m^-) := \sup_{\mathcal{U}_m} \mathbb{E}_{\mathcal{U}_m^{x, t_m^-}} \left[W_T + \gamma \left(w + \frac{1}{\alpha} \min(W_T - w, 0) \right) \right].$$

Terminal condition: denoting the terminal reward functional by $\Phi(x, w)$, we have

$$(3.8) \quad V(x, w, T^-) = \Phi(x, w) := (e^s + e^b) + \gamma \left(w + \frac{1}{\alpha} \min((e^s + e^b) - w, 0) \right), \quad x = (s, b).$$

Between $[t_m^+, t_{m+1}^-]$, $m = M-1, \dots, 0$, propagate via the uncontrolled kernel $g(\cdot; \Delta t)$:

$$(3.9) \quad V(x, w, t_m^+) = \int_{\mathbb{R}^2} V(y, w, t_{m+1}^-) g(y - x; \Delta t) dy.$$

At t_m , define the intervention operator for any test functional F and control $u \in \mathcal{Z}$,

$$(3.10) \quad (\mathcal{M}_u F)(x, w, t_m^-; q_m) := F(s^+(x, q_m, u), b^+(x, q_m, u), w, t_m^+),$$

where $(s^+(\cdot), b^+(\cdot))$ as in (3.3). The optimal control (for fixed w) is

$$(3.11) \quad u_m^*(\cdot; w) \in \arg \max_{u \in \mathcal{Z}} (\mathcal{M}_u V)(x, w, t_m^-; q_m).$$

The pre-intervention auxiliary value at t_m^- is

$$(3.12) \quad V(x, w, t_m^-) = (\mathcal{M}_{u_m^*(\cdot; w)} V)(x, w, t_m^-; q_m).$$

At inception, select the optimal threshold and value:

$$(3.13) \quad w^* = \arg \sup_{w \geq 0} V(x_0, w, t_0^-), \quad V(x_0, t_0^-) = V(x_0, w^*, t_0^-).$$

The pre-commitment optimal policy is $\mathcal{U}_0^* = \{u_0^*(\cdot; w^*), \dots, u_{M-1}^*(\cdot; w^*)\}$.

PROPOSITION 3.3 (Equivalence to the pre-commitment problem). *The lifted recursion (3.7)–(3.13) is equivalent to (3.6); see [25].*


3.3 Localization and problem statement Since w^* is finite by Lemma 3.2, we truncate the threshold domain to $\Gamma = [0, w_{\max}]$, where $w_{\max} > 0$ is chosen sufficiently large. We localize the (s, b) -domain to the rectangle

$$(3.14) \quad \Omega = [s_{\min}^\dagger, s_{\max}^\dagger] \times [b_{\min}^\dagger, b_{\max}^\dagger],$$

with $s_{\min}^\dagger < s_{\min} < 0 < s_{\max} < s_{\max}^\dagger$ and $b_{\min}^\dagger < b_{\min} < 0 < b_{\max} < b_{\max}^\dagger$ chosen so that boundary truncation errors are negligible (see, e.g. [33, 8]). On Ω_{out} , the terminal reward (3.8) suggests exponential leading-order forms (e^s , e^b , e^{s+b} , or a constant) depending on which boundary is approached. The next lemma shows that these forms propagate over one time step via explicit CF factors.

We partition Ω into the interior and four boundary sub-domains:

$$(3.15) \quad \begin{aligned} \Omega_{\text{in}} &= (s_{\min}, s_{\max}) \times (b_{\min}, b_{\max}), \\ \Omega_{s_{\max}} &= [s_{\max}, s_{\max}^{\dagger}] \times [b_{\min}^{\dagger}, b_{\max}], \\ \Omega_{s_{\max}}^{b_{\max}} &= [s_{\max}, s_{\max}^{\dagger}] \times [b_{\max}, b_{\max}^{\dagger}], \\ \Omega_{b_{\max}} &= [s_{\min}^{\dagger}, s_{\max}] \times [b_{\max}, b_{\max}^{\dagger}], \\ \Omega_{s_{\min}}^{b_{\min}} &= \Omega \setminus \Omega_{\text{in}} \setminus \Omega_{s_{\max}} \setminus \Omega_{b_{\max}} \setminus \Omega_{s_{\max}}^{b_{\max}}, \\ \Omega_{\text{out}} &:= \Omega \setminus \Omega_{\text{in}}. \end{aligned}$$



An illustration of the sub-domains for the localized problem at each t (fixed $w \in \Gamma$) is given in Figure 3.1.

Fig. 3.1: Spatial sub-domains at each t (fixed $w \in \Gamma$).

LEMMA 3.4 (One-step propagation for exponential functions). *Let $\phi : \mathbb{R}^2 \times [0, T] \rightarrow \mathbb{R}$. Under Assumption 2.1 (A5), for any $a \in \{(1, 0), (0, 1), (1, 1), (0, 0)\}$, if at t_{m+1}^- , $\phi(x, t_{m+1}^-) = A(t_{m+1}^-) e^{a \cdot x}$, where $A(t)$ denotes a generic unknown function of time, then applying the convolution integral (3.9) with ϕ in place of V yields*

$$\phi(x, t_m^+) = \int_{\mathbb{R}^2} \phi(y, t_{m+1}^-) g(y - x; \Delta t) dy = G(-i a; \Delta t) \phi(x, t_{m+1}^-),$$

where $G(-i a; \Delta t) \in (0, \infty)$ by Assumption 2.1 (A5).

Proof. A change of variables $y = x + z$ gives $\int_{\mathbb{R}^2} e^{a \cdot z} g(z; \Delta t) dz = G(-i a; \Delta t)$. \square

We encode the boundary type by the selector $a : \Omega_{\text{out}} \rightarrow \{(1, 0), (0, 1), (1, 1), (0, 0)\}$,

$$(3.16) \quad a(x) = (1, 0) \mathbf{1}_{\Omega_{s_{\max}}}(x) + (0, 1) \mathbf{1}_{\Omega_{b_{\max}}}(x) + (1, 1) \mathbf{1}_{\Omega_{s_{\max}}^{b_{\max}}}(x) + (0, 0) \mathbf{1}_{\Omega_{s_{\min}}^{b_{\min}}}(x).$$

Define the boundary operator

$$(3.17) \quad (\mathcal{B}V)(x, w, t_m^+) := G(-i a(x); \Delta t) V(x, w, t_{m+1}^-), \quad x \in \Omega_{\text{out}}.$$

No intervention is applied on Ω_{out} , so we also carry values from t_m^+ to t_m^- on the boundary. Hence, for $(x, t_m^\pm) \in \Omega_{\text{out}} \times \{t_m^\pm\}$,

$$(3.18) \quad V(x, w, t_m^-) = V(x, w, t_m^+) = (\mathcal{B}V)(x, w, t_m^+).$$

On the interior Ω_{in} , (3.9) is approximated by the truncated 2-D convolution

$$(3.19) \quad V(x, w, t_m^+) \simeq \int_{\Omega} V(y, w, t_{m+1}^-) g(y - x; \Delta t) dy, \quad x \in \Omega_{\text{in}},$$

where $x = (s, b)$, $y = (y_s, y_b)$, and $dy \equiv dy_s dy_b$.

DEFINITION 3.5 (Localized Mean-CVaR formulation). *The localized value function at time t_0^- is $V(x_0, w^*, t_0^-)$, where w^* is obtained by the outer search (3.13), and $\mathcal{U}_0^* = \{u_0^*(\cdot; w^*), \dots, u_{M-1}^*(\cdot; w^*)\}$ is the associated optimal control.*

For each fixed $w \in \Gamma$, the function $V(x, w, t)$ on $\Omega \times \mathcal{T}$ is specified by: (i) terminal condition (3.8); (ii) boundary propagation (3.18) on Ω_{out} for $m = M - 1, \dots, 0$; (iii) interior propagation (3.19) on Ω_{in} for $m = M - 1, \dots, 0$; and (iv) interior intervention (3.11)–(3.12) (with no intervention on Ω_{out}).

Remark 3.6 (Uniqueness and regularity of the localized inner problem). For each fixed $w \in \Gamma$ and each $m = 0, \dots, M$, the localized Bellman recursion admits a unique bounded solution on Ω ; moreover $V(\cdot, w, t_m^-)$ is continuous in (s, b) on Ω_{in} (cf. [33, Proposition 3.1]).

4 Fourier-trained transition kernel approximation For notational simplicity, we suppress the explicit dependence on Δt and write $g(\cdot) \equiv g(\cdot; \Delta t) : \mathbb{R}^2 \rightarrow \mathbb{R}$ for the (uncontrolled) transition density, with CF $G(\cdot) \equiv G(\cdot; \Delta t)$. We adopt the same convention for the kernel approximation $\widehat{g}(\cdot)$ and its CF $\widehat{G}(\cdot)$.

4.1 Description We approximate $g(\cdot)$ by a single-hidden-layer FFNN with Gaussian activation [11]. In 2D it is convenient to parameterize the network as a finite Gaussian mixture. Specifically, for $y = (y_s, y_b) \in \mathbb{R}^2$ we take

$$(4.1) \quad \phi(y; \mu, \Sigma) := \frac{1}{(2\pi)^{|\Sigma|^{1/2}}} \exp\left(-\frac{1}{2}(y-\mu)^\top \Sigma^{-1}(y-\mu)\right), \quad \mu \in \mathbb{R}^2, \quad \Sigma \in \mathbb{R}^{2 \times 2},$$

and define the class of single-layer FFNNs with Gaussian activation by

$$(4.2) \quad \Sigma(\phi) = \left\{ \widehat{g} : \mathbb{R}^2 \rightarrow \mathbb{R} \mid \widehat{g}(y; \theta) = \sum_{n=1}^N \beta_n \phi(y; \mu_n, \Sigma_n), \theta \in \Theta, N \in \mathbb{N} \right\},$$

where

$$(4.3) \quad \mu_n = (\mu_n^s, \mu_n^b) \in \mathbb{R}^2, \quad \Sigma_n = \begin{pmatrix} (\sigma_n^s)^2 & \rho_n \sigma_n^s \sigma_n^b \\ \rho_n \sigma_n^s \sigma_n^b & (\sigma_n^b)^2 \end{pmatrix},$$

with $\sigma_n^s > 0$, $\sigma_n^b > 0$, and $\rho_n \in (-1, 1)$, so each Σ_n is symmetric positive-definite.

Bounded parameter space and positivity. By Assumptions 2.1 (A1)–(A2), $g(\cdot; \Delta t) \in L_1(\mathbb{R}^2) \cap L_2(\mathbb{R}^2)$. The approximation results in [11, Cor. 3.2, Thm. 3.4] imply that single-hidden-layer Gaussian networks are L_2 -dense on $L_1 \cap L_2$ even with bounded parameters. We therefore restrict to the bounded parameter set

$$(4.4) \quad \Theta = \left\{ \theta = \{(\beta_n, \mu_n^s, \mu_n^b, \sigma_n^s, \sigma_n^b, \rho_n)\}_{n=1}^N : \beta_n \geq 0, \sum_{n=1}^N \beta_n = 1, \right. \\ \left. |\mu_n^s|, |\mu_n^b| \leq \bar{\mu}, \quad 0 < \sigma_{\min} \leq \sigma_n^s, \sigma_n^b \leq \sigma_{\max}, \quad |\rho_n| \leq \bar{\rho} < 1 \right\},$$

for fixed finite $\bar{\mu}, \sigma_{\max}$, positive σ_{\min} , and $0 < \bar{\rho} < 1$. Under (4.4), each Σ_n is uniformly positive-definite. Moreover, since $\phi(\cdot; \mu_n, \Sigma_n) \geq 0$ integrates to one and $\{\beta_n\}$ lies on the probability simplex, $\widehat{g}(\cdot; \theta)$ is a proper density: $\widehat{g} \geq 0$ and $\int_{\mathbb{R}^2} \widehat{g} = 1$.

Closed-form CF. Using the Fourier pair in (2.3), the CF of $\widehat{g}(\cdot)$ is available in closed form. For $\eta = (\eta_s, \eta_b) \in \mathbb{R}^2$,

$$(4.5) \quad \widehat{G}(\eta; \theta) = \int_{\mathbb{R}^2} e^{i\eta \cdot y} \widehat{g}(y; \theta) dy = \sum_{n=1}^N \beta_n \exp\left(i\eta^\top \mu_n - \frac{1}{2}\eta^\top \Sigma_n \eta\right).$$

L_2 -approximation error and Fourier-domain invariance. By Assumption 2.1 (A2), $g(\cdot) \in L_1(\mathbb{R}^2) \cap L_\infty(\mathbb{R}^2)$ and hence $g(\cdot) \in L_2(\mathbb{R}^2)$. The Fourier transform is an L_2 isometry up to a constant factor (see, e.g. [32]): if $f \in L_2(\mathbb{R}^2)$ with Fourier transform $\mathfrak{F}[f]$ under the convention (2.3), then

$$(4.6) \quad \int_{\mathbb{R}^2} |f(y)|^2 dy = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} |\mathfrak{F}[f](\eta)|^2 d\eta.$$

Combining (4.6) with [11, Thm. 3.4] yields the 2D approximation result below.

THEOREM 4.1 (*L_2 -approximation and Fourier-domain invariance; cf. [11], Thm. 3.4*).

For any $\varepsilon > 0$, there exists $\widehat{g}(\cdot; \theta_\varepsilon^) \in \Sigma(\phi)$ with $\theta_\varepsilon^* \in \Theta$ such that*

$$(4.7) \quad \int_{\mathbb{R}^2} |g(y) - \widehat{g}(y; \theta_\varepsilon^*)|^2 dy = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} |G(\eta) - \widehat{G}(\eta; \theta_\varepsilon^*)|^2 d\eta < \varepsilon,$$

where $\widehat{G}(\eta; \theta)$ is given in (4.5).

Proof. Existence follows by a 2D adaptation of [11, Thm. 3.4] for $\Sigma(\phi)$ with bounded parameters Θ . The equality in (4.7) follows from (4.6) applied to $f = g - \widehat{g}$. \square

4.2 Training loss and regularization We truncate the Fourier domain from \mathbb{R}^2 to $D_\eta = [-\eta', \eta']^2$, with $\eta' > 0$ chosen sufficiently large, and discretize D_η by a (possibly non-uniform) set of nodes $\{\eta_p\}_{p=1}^P \subset D_\eta$, $\eta_p = (\eta_p^s, \eta_p^b)$. Define

$$\delta_{\min} := \min_{p \neq q} \|\eta_p - \eta_q\|_2, \quad \delta_{\max} := \max_{1 \leq p \leq P} \min_{q \neq p} \|\eta_p - \eta_q\|_2,$$

and assume the sampling set is quasi-uniform: there exist constants $C_0, C_1 > 0$, independent of P and N , such that

$$(4.8) \quad C_0 P^{-1/2} \leq \delta_{\min} \leq \delta_{\max} \leq C_1 P^{-1/2}.$$

Let $\widehat{\Theta} \subset \Theta$ be the empirical parameter set. Using the closed-form target CF $G(\cdot)$ and the network CF $\widehat{G}(\cdot; \theta)$ from (4.5), we compute θ by minimizing the empirical loss

$$(4.9) \quad \text{Loss}_P(\theta) = \frac{1}{P} \sum_{p=1}^P |G(\eta_p) - \widehat{G}(\eta_p; \theta)|^2 + R_P(\theta), \quad \theta \in \widehat{\Theta}.$$

Here, $R_P(\theta)$ is the MAE regularization term

$$R_P(\theta) = \frac{1}{P} \sum_{p=1}^P \left(|\text{Re}_G(\eta_p) - \text{Re}_{\widehat{G}}(\eta_p; \theta)| + |\text{Im}_G(\eta_p) - \text{Im}_{\widehat{G}}(\eta_p; \theta)| \right),$$

where Re_G, Im_G denote the real and imaginary parts of G (and likewise for \widehat{G}). The empirical minimizer is

$$(4.10) \quad \widehat{\theta}^* = \arg \min_{\theta \in \widehat{\Theta}} \text{Loss}_P(\theta).$$

The MAE term concentrates accuracy in high-impact regions of G (e.g. near peaks and rapid variations) [11]. In practice, the sampling set $\{\eta_p\}$ is chosen to be denser in such regions.

4.3 Error decomposition We decompose the learned-kernel error into: (i) Fourier truncation on $\mathbb{R}^2 \setminus D_\eta$, (ii) empirical training error on the sampled set, and (iii) sampling (quadrature) error on D_η . The next lemma provides explicit bounds for $g - \widehat{g}$.

LEMMA 4.2 (2D training error decomposition). *Let $D_\eta = [-\eta', \eta']^2$ and let $\{\eta_p\}_{p=1}^P \subset D_\eta$ satisfy (4.8). Let $\widehat{\theta}^* \in \widehat{\Theta}$ be the empirical minimizer in (4.10). Assume that, for some $\varepsilon_1 > 0$ and $f \in \{G(\cdot), \widehat{G}(\cdot; \widehat{\theta}^*)\}$,*

$$(4.11) \quad \int_{\mathbb{R}^2 \setminus D_\eta} (|\text{Re}_f(\eta)| + |\text{Im}_f(\eta)|) d\eta < \varepsilon_1, \quad \int_{\mathbb{R}^2 \setminus D_\eta} |f(\eta)|^2 d\eta < \varepsilon_1,$$

and that $\text{Loss}_P(\widehat{\theta}^*) < \varepsilon_2$ and $R_P(\widehat{\theta}^*) < \varepsilon_3$. Moreover, suppose

$$(4.12) \quad C' := \sup_{\eta \in D_\eta, \theta \in \Theta} \|\nabla_\eta |G(\eta) - \widehat{G}(\eta; \theta)|^2\|_2 < \infty.$$

Then:

(i) Global L_2 -error bound.

$$(4.13) \quad \int_{\mathbb{R}^2} |g(y) - \widehat{g}(y; \widehat{\theta}^*)|^2 dy \leq \frac{1}{(2\pi)^2} \left\{ 4\varepsilon_1 + C_1 \varepsilon_2 + \frac{C' C_1^3}{2 P^{1/2}} \right\}.$$

(ii) Pointwise bound. For every $y \in \mathbb{R}^2$,

$$(4.14) \quad |g(y) - \widehat{g}(y; \widehat{\theta}^*)| \leq \frac{1}{(2\pi)^2} \left\{ 2\varepsilon_1 + C_1 \varepsilon_3 + \frac{C' C_1^3}{2 P^{1/2}} \right\}.$$

A proof of Lemma 4.2 is given in Appendix B.

We now show that, under Assumption 2.1 (A4) and the bounded parameter set (4.4), the Fourier mass outside the truncated domain $D_\eta = [-\eta', \eta']^2$ (and hence the Fourier truncation error) decays exponentially for both G and \widehat{G} .

LEMMA 4.3 (Fourier–domain truncation error bounds). *Suppose Assumption 2.1 (A4) holds for $G(\cdot)$. Recall $\widehat{G}(\cdot; \theta)$ from (4.5) with parameters bounded as in (4.4). There exist positive constants $C_G^{(p)}, c_G^{(p)}$, and, for each fixed finite N , constants $C_{\widehat{G}}^{(p)}, c_{\widehat{G}}^{(p)} > 0$ (uniform over all θ in the bounded parameter set with N terms) such that, for any $\eta' \geq R_{\text{tail}}$ and $p \in \{1, 2\}$,*

$$\int_{\mathbb{R}^2 \setminus [-\eta', \eta']^2} |G(\eta)|^p d\eta \leq C_G^{(p)} e^{-c_G^{(p)} (\eta')^\alpha}, \quad \int_{\mathbb{R}^2 \setminus [-\eta', \eta']^2} |\widehat{G}(\eta; \theta)|^p d\eta \leq C_{\widehat{G}}^{(p)} e^{-c_{\widehat{G}}^{(p)} (\eta')^2}.$$

A proof of Lemma 4.3 is given in Appendix C.

Let $h \in (0, 1)$ denote a target tolerance. The next corollary gives a sufficient choice of $\eta' = \eta'(h)$ so that the Fourier–domain truncation error for both G and \widehat{G} satisfies $\leq h^{1+\kappa}$ for any $\kappa \in (0, 1]$.

COROLLARY 4.4 (Fourier-domain truncation error $O(h^{1+\kappa})$). *Fix $p \in \{1, 2\}$ and $\kappa \in (0, 1]$. Under Assumption 2.1 (A4) and Lemma 4.3, define for $h \in (0, 1)$*

$$\eta'(h) := \max \left\{ R_{\text{tail}}, \left(\frac{1}{c_G^{(p)}} \log \frac{2C_G^{(p)}}{h^{1+\kappa}} \right)^{1/\alpha}, \left(\frac{1}{c_{\widehat{G}}^{(p)}} \log \frac{2C_{\widehat{G}}^{(p)}}{h^{1+\kappa}} \right)^{1/2} \right\}.$$

Then $\max_{f \in \{G, \widehat{G}\}} \int_{\mathbb{R}^2 \setminus [-\eta'(h), \eta'(h)]^2} |f(\eta)|^p d\eta \leq h^{1+\kappa}$.

To establish our error bounds in terms of h , we introduce the following sampling and training assumptions, indexed by a target tolerance $h \in (0, 1]$.

ASSUMPTION 4.5 (Sampling and training assumptions). *Let $h \in (0, 1)$ be a training tolerance parameter. Fix $\kappa \in (0, 1)$ and choose $\eta'(h)$ as in Corollary 4.4 so that, for $p \in \{1, 2\}$,*

$$\max_{f \in \{G, \widehat{G}\}} \int_{\mathbb{R}^2 \setminus [-\eta', \eta']^2} |f(\eta)|^p d\eta \leq c_1 h^{1+\kappa}, \quad \eta' = \eta'(h).$$

Assume the quasi-uniformity constant C_1 in (4.8) and the gradient bound C' in (4.12) are finite and independent of h . Assume further that the trained network satisfies: (i) $\text{Loss}_P(\widehat{\theta}^*) \leq c_2 h^{1+\kappa}$, (ii) $R_P(\widehat{\theta}^*) \leq c_3 h^{1+\kappa}$, and (iii) $P \geq c_P h^{-2(1+\kappa)}$, where $c_1, c_2, c_3, c_P > 0$ are constants independent of h .

COROLLARY 4.6 (h -dependent convergence rates). *Under Assumption 4.5, there exists a generic constant $C > 0$, independent of h (its value may change from line to line), such that:*

- (i) Global L_2 -error bound. $\|g - \widehat{g}\|_{L_2(\mathbb{R}^2)} \leq C h^{\frac{1+\kappa}{2}}$.
- (ii) Pointwise bound. For every $y \in \mathbb{R}^2$,

$$(4.15) \quad |g(y) - \widehat{g}(y; \widehat{\theta}^*)| \leq C h^{1+\kappa}.$$

A proof of Corollary 4.6 is given in Appendix D.

Remark 4.7 (Model capacity is implicit). Assumption 4.5 does not prescribe a specific growth law $N(h)$ for the network size. We only require that (for each h) there exist trained parameters (possibly with $N = N(h)$) that achieve the stated empirical error bounds. Thus, any architecture-size schedule that delivers $\text{Loss}_P(\widehat{\theta}^*) = O(h^{1+\kappa})$ and $R_P(\widehat{\theta}^*) = O(h^{1+\kappa})$ is admissible, and we keep $N(h)$ implicit.

5 Numerical methods

5.1 Discretization To compute the interior convolution on Ω_{in} and the boundary propagation on $\Omega \setminus \Omega_{\text{in}}$ on a single tensor grid, we pad each interior interval $[z_{\min}, z_{\max}]$ ($z \in \{s, b\}$) symmetrically by half its width (chosen sufficiently large so that boundary truncation errors are negligible; cf. Section 3.3): for $z \in \{s, b\}$,

$$(5.1) \quad z_{\min}^{\dagger} = z_{\min} - \frac{1}{2}(z_{\max} - z_{\min}), \quad z_{\max}^{\dagger} = z_{\max} + \frac{1}{2}(z_{\max} - z_{\min}).$$

Let $D_z := z_{\max} - z_{\min}$ and $D_z^{\dagger} := z_{\max}^{\dagger} - z_{\min}^{\dagger} = 2D_z$. For $z \in \{s, b\}$, let N_z be the number of intervals on $[z_{\min}, z_{\max}]$ and set $N_z^{\dagger} := 2N_z$ on $[z_{\min}^{\dagger}, z_{\max}^{\dagger}]$, so that

$$(5.2) \quad \Delta z = \frac{D_z}{N_z} = \frac{D_z^{\dagger}}{N_z^{\dagger}}, \quad z \in \{s, b\}.$$

Assume N_s and N_b are even and define the grid nodes by

$$(5.3) \quad s_k = \hat{s}_0 + k \Delta s, \quad k = -\frac{N_s^{\dagger}}{2}, \dots, \frac{N_s^{\dagger}}{2}, \quad b_j = \hat{b}_0 + j \Delta b, \quad j = -\frac{N_b^{\dagger}}{2}, \dots, \frac{N_b^{\dagger}}{2},$$

where $\hat{s}_0 := \frac{1}{2}(s_{\min}^{\dagger} + s_{\max}^{\dagger})$ and $\hat{b}_0 := \frac{1}{2}(b_{\min}^{\dagger} + b_{\max}^{\dagger})$. For $z \in \{s, b\}$, define the global and interior index sets

$$\mathbb{N}_z := \left\{ -\frac{N_z}{2} + 1, \dots, \frac{N_z}{2} - 1 \right\} \subset \mathbb{N}_z^{\dagger} := \left\{ -\frac{N_z^{\dagger}}{2}, \dots, \frac{N_z^{\dagger}}{2} \right\},$$

so both the interior and boundary updates use the same global node set $\{(s_k, b_j)\}$.

We discretize the admissible control actions $\mathcal{Z} = [0, 1]$ using N_u uniform intervals, yielding nodes $\{u_l\}_{l=0}^{N_u}$, and discretize the threshold domain $\Gamma = [0, w_{\max}]$ at uniform nodes $\{w_c\}_{c=0}^{N_w}$, with spacings Δu and Δw .

Single refinement parameter. As is common in the literature (e.g., [21, 6][Eq. (4.1)], [10][Eq. (3.24)]), we introduce a single parameter $h \in (0, 1)$ to index both mesh refinement and the training tolerance from Assumption 4.5. For the spatial, control, and threshold grids we set

$$(5.4) \quad \Delta s = C_1 h, \quad \Delta b = C_2 h, \quad \Delta u = C_3 h, \quad \Delta w = C_4 h,$$

with $C_1, C_2, C_3, C_4 > 0$ independent of h . We also index the Fourier truncation and sampling budget by h via $\eta' = \eta'(h)$ (Corollary 4.4) and $P = P(h)$ with $P(h) \geq c_P h^{-2(1+\kappa)}$ (Assumption 4.5).

For each discretized threshold value w_c , let $V_{k,j,c}^{m,\circ}$ be a numerical approximation to $V(s_k, b_j, w_c, t_m^{\circ})$ at the node $(s_k, b_j, w_c, t_m^{\circ})$, where $t_m^{\circ} \in \{t_m, t_m^{\pm}\}$. Given nodal values $\{V_{l,d,c}^{m,+}\}$ on the global grid, we denote by $\mathcal{I}[\{V_c^{m,+}\}](s, b)$ the bilinear interpolant on Ω (for this fixed w_c).

Because the convolution integral (3.19) evaluates the kernel at $(y - x)$, we use the index-difference notation

$$(5.5) \quad g_{l-k, d-j} := g((s_l - s_k, b_d - b_j); \Delta t), \quad \hat{g}_{l-k, d-j} := \hat{g}((s_l - s_k, b_d - b_j); \hat{\theta}^*),$$

for $(k, j) \in \mathbb{N}_s \times \mathbb{N}_b$ and $(l, d) \in \mathbb{N}_s^{\dagger} \times \mathbb{N}_b^{\dagger}$, where $\hat{\theta}^*$ is the minimizer (4.10). Since $\hat{g}(\cdot; \hat{\theta}^*)$ is available in closed form on \mathbb{R}^2 (Section 4), these evaluations are well-defined.

5.2 Numerical scheme We compute the lifted value function backward in time for each fixed discretized threshold $w_c \in \{w_c\}_{c=0}^{N_w}$.

Terminal condition. At maturity $t_M = T$, we implement (3.8) on Ω via

$$(5.6) \quad V_{k,j,c}^{M,-} = \Phi(s_k, b_j, w_c), \quad (k, j) \in \mathbb{N}_s^{\dagger} \times \mathbb{N}_b^{\dagger}.$$

Boundary sub-domains. For nodes (s_k, b_j, t_m^\pm) on $\Omega_{\text{out}} \times \{t_m^\pm\}$, $m = M-1, \dots, 0$, we enforce the asymptotic boundary conditions (3.18) as follows:

$$(5.7) \quad V_{k,j,c}^{m,-} = V_{k,j,c}^{m,+} = G(-i a(s_k, b_j); \Delta t) V_{k,j,c}^{m+1,-}.$$

Interior sub-domain. For nodes $(s_k, b_j, t_m^+) \in \Omega_{\text{in}} \times \{t_m^+\}$, $m = M-1, \dots, 0$, the truncated 2-D convolution integral (3.19) is approximated by

$$(5.8) \quad V_{k,j,c}^{m,+} = \Delta s \Delta b \sum_{l \in \mathbb{N}_s^\dagger} \sum_{d \in \mathbb{N}_b^\dagger} \varphi_{l,d} \widehat{g}_{l-k, d-j} V_{l,d,c}^{m+1,-}, \quad (k, j) \in \mathbb{N}_s \times \mathbb{N}_b.$$

Here, $\{\varphi_{l,d}\}$ are the composite trapezoidal weights (unit weight in the interior, $\frac{1}{2}$ on edges, and $\frac{1}{4}$ at corners). The kernel weights $\widehat{g}_{l-k, d-j}$ are evaluated directly from the learned Gaussian-mixture representation in Section 4 (via (4.1) and (4.2) with θ^*).

Intervention (interior only). For each interior node $(s_k, b_j) \in \Omega_{\text{in}}$ at decision time t_m , $m = M-1, \dots, 0$, we enforce (3.11)–(3.12) on the discrete control set $\{u_\iota\}_{\iota=0}^{N_u}$:

$$(5.9) \quad \begin{aligned} V_{k,j,c}^{m,-} &= \max_{u_\iota \in \{u_\iota\}_{\iota=0}^{N_u}} \mathcal{I}[\{V_c^{m,+}\}](s_{k,j}^{m,+}(u_\iota), b_{k,j}^{m,+}(u_\iota)), \quad (k, j) \in \mathbb{N}_s \times \mathbb{N}_b, \\ s_{k,j}^{m,+}(u) &= s^+(s_k, b_j, q_m, u), \quad b_{k,j}^{m,+}(u) = b^+(s_k, b_j, q_m, u), \end{aligned}$$

where $(s^+(\cdot), b^+(\cdot))$ are given by (3.3) and $\mathcal{I}[\cdot]$ denotes bilinear interpolation. This step yields the numerically computed optimal control $u_{k,j,c}^{m,*} \equiv u_{k,j}^{m,*}(w_c)$:

$$(5.10) \quad u_{k,j,c}^{m,*} \equiv u_{k,j}^{m,*}(w_c) \in \arg \max_{u_\iota \in \{u_\iota\}_{\iota=0}^{N_u}} \mathcal{I}[\{V_c^{m,+}\}](s_{k,j}^{m,+}(u_\iota), b_{k,j}^{m,+}(u_\iota)).$$

Initial time t_0 . At t_0 , the post-intervention values $\{V_c^{0,-}\}$ are available on the grid. Define $\widehat{V}(s, b, w_c, t_0^-) := \mathcal{I}[\{V_c^{0,-}\}](s, b)$ and determine the pre-commitment threshold and value by searching over $\{w_c\}_{c=0}^{N_w}$:

$$(5.11) \quad w_{c*} \in \arg \max_{\{w_c\}_{c=0}^{N_w}} \widehat{V}(s_0, b_0, w_c, t_0^-), \quad V_h(s_0, b_0, t_0^-) := \widehat{V}(s_0, b_0, w_{c*}, t_0^-).$$

Here, $V_h(s_0, b_0, t_0^-)$ is the numerical approximation (at refinement parameter h) to $V(s_0, b_0, t_0^-)$; the corresponding optimal controls $u_{k,j,c*}^{m,*}$ are obtained and stored during the backward recursion.

5.3 Efficient implementation We accelerate the 2-D discrete convolution (5.8) using FFTs. Since $\widehat{g}_{l-k, d-j}$ depends only on index differences, the discrete convolution operator is Toeplitz-block-Toeplitz and can be embedded into a 2-D circular convolution on an augmented grid [10, 34]. Concretely, with $*$ denoting circular convolution, (5.8) can be written in the circular-convolution form

$$(5.12) \quad \mathbf{V}_c^{m,+} = \Delta s \Delta b \widehat{\mathbf{g}} * \mathbf{V}_c^{m+1,-}.$$

Here, $\widehat{\mathbf{g}}$ and $\mathbf{V}_c^{m+1,-}$ respectively denote the appropriately augmented (zero-padded and re-indexed) kernel and value arrays of size $(3N_s - 1) \times (3N_b - 1)$ constructed from $\{\widehat{g}_{l-k, d-j}\}$ and $\{\varphi_{l,d} V_{l,d,c}^{m+1,-}\}_{l \in \mathbb{N}_s^\dagger, d \in \mathbb{N}_b^\dagger}$ (see [10, 34]). The notation $*$ denotes circular convolution. The circular convolution (5.12) is then computed via FFT/iFFT:

$$(5.13) \quad \mathbf{V}_c^{m,+} = \Delta s \Delta b \text{iFFT} \{ \text{FFT} \{ \mathbf{V}_c^{m+1,-} \} \circ \text{FFT} \{ \widehat{\mathbf{g}} \} \}.$$

After (5.13), we extract the components corresponding to $(k, j) \in \mathbb{N}_s \times \mathbb{N}_b$ to obtain the interior values $V_{k,j,c}^{m,+}$ on $\Omega_{\text{in}} \times \{t_m^+\}$.

6 Convergence analysis In this section, we use a Barles–Souganidis–type analysis [2] to prove convergence, as $h \rightarrow 0$, to the localized formulation in Definition 3.5. Fix a discrete threshold node $w_c \in \Gamma_h$. We verify ℓ_∞ -stability, monotonicity, and pointwise consistency for the associated inner recursion, which applies (5.6) on $\Omega \times \{T\}$, (5.7) on Ω_{out} (updated at t_m^+ and carried to t_m^-), and (5.8)–(5.9) on $\Omega_{\text{in}} \times \{t_m^\pm\}$, for $m = M - 1, \dots, 0$.

For an arbitrary $\theta \in \widehat{\Theta}$, write $\widehat{g}^{(\theta)}(y) := \widehat{g}(y; \theta)$ and $\widehat{g}_{l-k, d-j}^{(\theta)} := \widehat{g}((s_l - s_k, b_d - b_j); \theta)$. When $\theta = \widehat{\theta}^*$, we drop $\widehat{\theta}^*$ and revert to the convention adopted in (5.5).

A *discrete kernel-mass bound*. Assumption 2.1 (A4) implies $g(\cdot; \Delta t) \in C^\infty(\mathbb{R}^2)$ with bounded derivatives [30, Ch. 3]. In particular, for any multi-index β ,

$$\begin{aligned} \partial^\beta g(y; \Delta t) &= \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} (-i\eta)^\beta e^{-i\eta \cdot y} G(\eta; \Delta t) d\eta, \\ \|\partial^\beta g(\cdot; \Delta t)\|_\infty &\leq \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} \|\eta\|_2^{|\beta|} |G(\eta; \Delta t)| d\eta < \infty. \end{aligned}$$

For $(s_k, b_j) \in \Omega_{\text{in}}$, we write $x_{k,j} = (s_k, b_j)$, and for nodes on Ω , we write $y_{l,d} = (s_l, b_d)$.

Define the (bounded) extension rectangle

$$(6.1) \quad \Omega_{\text{ext}} := [s_{\min}^\dagger - s_{\max}, s_{\max}^\dagger - s_{\min}] \times [b_{\min}^\dagger - b_{\max}, b_{\max}^\dagger - b_{\min}],$$

so that $\Omega - \{x_{k,j}\} \subseteq \Omega_{\text{ext}} \forall x_{k,j} \in \Omega_{\text{in}}$. Then for any $(s_k, b_j) \in \Omega_{\text{in}}$ and any $\theta \in \widehat{\Theta}$,

$$\begin{aligned} \Delta s \Delta b \sum_{l,d} \varphi_{l,d} |\widehat{g}_{l-k, d-j}^{(\theta)}| &= \Delta s \Delta b \sum_{(y_{l,d} \in \Omega)} \varphi_{l,d} \left| \widehat{g}^{(\theta)}(y_{l,d} - x_{k,j}) \right| \\ &\leq \Delta s \Delta b \sum_{(y_{l,d} \in \Omega)} \varphi_{l,d} \left(g(y_{l,d} - x_{k,j}) + |\widehat{g}^{(\theta)} - g|(y_{l,d} - x_{k,j}) \right) \\ &\leq \int_{\Omega - \{x_{k,j}\}} g(z; \Delta t) dz + C_q h^2 + |\Omega| \|\widehat{g}^{(\theta)} - g\|_{L_\infty(\Omega_{\text{ext}})} \\ (6.2) \quad &\leq 1 + C_q h^2 + |\Omega| \|\widehat{g}^{(\theta)} - g\|_{L_\infty(\Omega_{\text{ext}})}. \end{aligned}$$

Here, $C_q > 0$ is independent of h and (k, j) (composite trapezoid error on a uniform grid).

Now, in (6.2), set $\theta = \widehat{\theta}^*$. By Corollary 4.6 (under Assumption 4.5), we have $\|\widehat{g} - g\|_{L_\infty(\mathbb{R}^2)} \leq C'' h^{1+\kappa}$. Hence, for sufficiently small h ,

$$(6.3) \quad \Delta s \Delta b \sum_{l,d} |\widehat{g}_{l-k, d-j}| \leq 1 + C' h^{1+\kappa} + C_q h^2 \leq 1 + \varepsilon(h) \leq e^{\varepsilon(h)}, \quad \varepsilon(h) = C h^{1+\kappa}.$$

6.1 Stability We now show ℓ_∞ -stability for fixed w_c .

LEMMA 6.1 (ℓ_∞ -stability for fixed w_c). *Fix $w_c \in \Gamma_h$ and let $h > 0$ be the global refinement parameter with (5.4) and Assumption 4.5 satisfied. If bilinear interpolation is used in (5.9), then the scheme (5.6), (5.7), and (5.8)–(5.9) is ℓ_∞ -stable for each fixed w_c : there exist constants $h_0 > 0$ and $0 < C < \infty$, independent of h , such that for all $0 < h \leq h_0$ and all $m = 0, \dots, M$,*

$$\|V_c^{m,-}\|_\infty \leq C, \quad \text{where} \quad \|V_c^{m,-}\|_\infty := \max_{k \in \mathbb{N}_b^{\dagger}, j \in \mathbb{N}_b^{\dagger}} |V_{k,j,c}^{m,-}|.$$

A proof of Lemma 6.1 is given in Appendix E.

6.2 Consistency We recall $s^+(x, q_m, u)$ and $b^+(x, q_m, u)$ from (3.3), and write $x^+(x, q_m, u) := (s^+(x, q_m, u), b^+(x, q_m, u))$. Denote $\hat{x} = (x, w)$ and $\hat{x}^{m, \circ} = (x, w, t_m^\circ)$ with $t_m^\circ \in \{t_m, t_m^\pm\}$. We write the localized backward recursion at the reference point $\hat{x}^{m, -}$ via the operator $\mathcal{D}(\cdot)$: $V(\hat{x}^{m, -}) = \mathcal{D}(\hat{x}^{m, -}, V^{m+1, -}) = \dots$

$$\dots = \begin{cases} \sup_{u \in \mathcal{Z}} \int_{\Omega} V(y, w, t_{m+1}^-) g(y - x^+(x, q_m, u); \Delta t) dy, & x \in \Omega_{\text{in}}, \quad m = M-1, \dots, 0, \\ G(-i a(x); \Delta t) V(x, w, t_{m+1}^-), & x \in \Omega_{\text{out}}, \quad m = M-1, \dots, 0, \\ \Phi(x, w), & x \in \Omega, \quad m = M. \end{cases}$$

Let $\Omega^h \times \Gamma^h$ be the computational grid, with Ω_{in}^h and Ω_{out}^h the interior/boundary sub-grids. The numerical scheme at the reference node $\hat{x}_{k,j,c}^{m,-} = (s_k, b_j, w_c, t_m^-)$ is written via the discrete operator $\mathcal{D}_h(\cdot)$: $V_{k,j,c}^{m,-} = \mathcal{D}_h(\hat{x}_{k,j,c}^{m,-}, \{V_{l,d,c}^{m+1,-}\}) = \dots$

$$\dots = \begin{cases} \max_{u \in \mathcal{Z}_h} \mathcal{I}[\{V_c^{m,+}\}](s_{k,j}^{m,+}(u), b_{k,j}^{m,+}(u)) & x_{k,j} \in \Omega_{\text{in}}^h, \quad m = M-1, \dots, 0, \\ G(-i a(s_k, b_j); \Delta t) V_{k,j,c}^{m+1,-}, & x_{k,j} \in \Omega_{\text{out}}^h, \quad m = M-1, \dots, 0, \\ \Phi(s_k, b_j, w_c), & x_{k,j} \in \Omega^h, \quad m = M. \end{cases}$$

where $(s_{k,j}^{m,+}(u), b_{k,j}^{m,+}(u)) = (s^+(s_k, b_j, q_m, u), b^+(s_k, b_j, q_m, u))$ given as in (3.3).

LEMMA 6.2 (Pointwise consistency; fixed w_c). *Fix $w_c \in \Gamma_h$ and let $h > 0$ satisfy (5.4) and Assumption 4.5. For any smooth test function $\phi(\cdot, w_c, \cdot) \in C^\infty(\Omega \times [0, T])$, denote $\phi^{m+1,-}(y) := \phi(y, w_c, t_{m+1}^-)$. Then, for all $m = M-1, \dots, 0$ and all $x_{k,j} \in \Omega_{\text{in}}^h$,*

$$\mathcal{D}_h(\hat{x}_{k,j,c}^{m,-}, \{\phi(\hat{x}_{l,d,c}^{m+1,-})\}_{l,d}) - \mathcal{D}(\hat{x}_{k,j,c}^{m,-}, \phi^{m+1,-}) = \mathcal{O}(h + h^{1+\kappa} + h^2).$$

Moreover, $\mathcal{D}_h(\cdot) = \mathcal{D}(\cdot)$ on Ω^h for $m = M$, and on Ω_{out}^h for $m = M-1, \dots, 0$.

A proof of Lemma 6.2 is given in Appendix F.

6.3 Monotonicity

LEMMA 6.3 (Monotonicity). *For fixed w_c and any bounded data sets $\{v_{l,d,c}^{m+1,-}\}_{l,d}$ and $\{z_{l,d,c}^{m+1,-}\}_{l,d}$ with $\{v_{l,d,c}^{m+1,-}\}_{l,d} \leq \{z_{l,d,c}^{m+1,-}\}_{l,d}$ (componentwise), the discrete operator satisfies*

$$\mathcal{D}_h(\hat{x}_{k,j,c}^{m,-}, \{v_{l,d,c}^{m+1,-}\}_{l,d}) \leq \mathcal{D}_h(\hat{x}_{k,j,c}^{m,-}, \{z_{l,d,c}^{m+1,-}\}_{l,d}), \quad \forall (k, j), \quad m = M-1, \dots, 0.$$

Proof. On Ω_{out}^h , the update is multiplication by $G(-i a(\cdot); \Delta t) > 0$, so order is preserved. On Ω_{in}^h , define the post-propagation arrays

$$\tilde{v}_{k,j} = \Delta s \Delta b \sum_{l,d} \varphi_{l,d} \hat{g}_{l-k, d-j} v_{l,d,c}^{m+1,-}, \quad \tilde{z}_{k,j} = \Delta s \Delta b \sum_{l,d} \varphi_{l,d} \hat{g}_{l-k, d-j} z_{l,d,c}^{m+1,-}.$$

Since $\varphi_{l,d} \geq 0$ and $\hat{g}_{l-k, d-j} \geq 0$ (mixture weights $\beta_n \geq 0$), we have $\tilde{v}_{k,j} \leq \tilde{z}_{k,j}$ componentwise on the full grid (including the boundary update). Bilinear interpolation uses nonnegative weights summing to one, hence preserves order:

$$\mathcal{I}[\{\tilde{v}\}](x_{k,j}^{m,+}(u)) \leq \mathcal{I}[\{\tilde{z}\}](x_{k,j}^{m,+}(u)), \quad \forall u \in \mathcal{Z}_h.$$

Taking $\max_{u \in \mathcal{Z}_h}$ preserves the inequality, yielding the claim. \square

A generic grid point in Ω_{in}^h is denoted by (s_h, b_h) , and we write $w_h \in \Gamma^h$ for a discrete threshold node. Let V_h denote the numerical solution produced by (5.6)–(5.8) (with bilinear interpolation in (5.9)), and recall V denotes the value function of the localized continuous problem.

LEMMA 6.4 (Convergence of the inner problem; fixed discrete threshold). *Fix the global refinement parameter $h > 0$ with (5.4) and Assumption 4.5 satisfied, and fix any $w_h \in \Gamma^h$. Let $(s', b') \in \Omega_{in}$ and let $\{(s_h, b_h)\}_{h \downarrow 0}$ be any sequence with $(s_h, b_h) \in \Omega_{in}^h$ and $(s_h, b_h) \rightarrow (s', b')$ as $h \rightarrow 0$. Then, for each $m \in \{M-1, \dots, 0\}$,*

$$(6.4) \quad |V_h(s_h, b_h, w_h, t_m^-) - V(s', b', w_h, t_m^-)| \longrightarrow 0 \quad \text{as } h \rightarrow 0.$$

Proof. Fix $w_h \in \Gamma^h$. By Lemma 6.1, Lemma 6.3, and Lemma 6.2, the scheme is ℓ_∞ -stable, monotone, and pointwise consistent (with error $\mathcal{O}(h) + \mathcal{O}(h^{1+\kappa}) + \mathcal{O}(h^2)$). Thus the Barles–Souganidis half-relaxed limits argument applies [2]. By the uniqueness of the localized inner value function for fixed threshold (Remark 3.6), the upper and lower limits coincide with $V(\cdot, w_h, t_m^-)$, hence the desired result (6.4). \square

LEMMA 6.5 (Convergence of the outer optimization). *Fix $x_0 \in \Omega$ and, for each $h > 0$, choose $x_0^h \in \Omega^h$ with $x_0^h \rightarrow x_0$. Define the outer objectives at t_0^- by*

$$F(w) := V(x_0, w, t_0^-), \quad F_h(w_h) := V_h(x_0^h, w_h, t_0^-),$$

for $w \in \Gamma$ and $w_h \in \Gamma_h$. Assume (5.4) and Assumption 4.5. Then:

(i) Convergence of optimal values.

$$\left| \sup_{w \in \Gamma} F(w) - \sup_{w_h \in \Gamma_h} F_h(w_h) \right| \leq \mathcal{O}(h) + \mathcal{O}(h^{1+\kappa}) + \mathcal{O}(h^2) + L_w \Delta w \xrightarrow{h \rightarrow 0} 0,$$

where $L_w := \gamma(1 + 1/\alpha)$.

(ii) Convergence of maximizers. *Let $w_h^* \in \arg \max_{w_h \in \Gamma_h} F_h(w_h)$. Every accumulation point \bar{w} of $\{w_h^*\}$ belongs to $\arg \max_{w \in \Gamma} F(w)$. If the maximizer of F on Γ is unique, then $w_h^* \rightarrow w^*$ as $h \rightarrow 0$.*

Proof. (i) Let $\varepsilon_h := \sup_{w_h \in \Gamma_h} |F_h(w_h) - F(w_h)|$. By Lemma 6.4 (at $m = 0$) and the fact that stability/consistency constants do not depend on w (the threshold enters only through the terminal objective functional, which is matched exactly), $\varepsilon_h = \mathcal{O}(h) + \mathcal{O}(h^{1+\kappa}) + \mathcal{O}(h^2)$. Moreover, for $w_1, w_2 \in \Gamma$,

$$|F(w_1) - F(w_2)| \leq \gamma \left(1 + \frac{1}{\alpha}\right) |w_1 - w_2| =: L_w |w_1 - w_2|,$$

so F is L_w -Lipschitz on Γ . The standard grid-approximation argument then yields the stated bound.

(ii) Let $w_h^* \in \arg \max_{w_h \in \Gamma_h} F_h(w_h)$ and let $w^* \in \arg \max_{w \in \Gamma} F(w)$. Along any subsequence with $w_h^* \rightarrow \bar{w} \in \Gamma$, we have $F(\bar{w}) \geq F(w_h^*) - L_w |w_h^* - \bar{w}| \geq F_h(w_h^*) - \varepsilon_h - L_w |w_h^* - \bar{w}|$, and since $F_h(w_h^*) = \sup_{\Gamma_h} F_h$, $F_h(w_h^*) = \sup_{\Gamma_h} F_h \geq \sup_{\Gamma} F - \varepsilon_h - L_w \Delta w$, it follows that

$$F(\bar{w}) \geq \sup_{\Gamma} F - 2\varepsilon_h - L_w \Delta w - L_w |w_h^* - \bar{w}|.$$

Letting $h \rightarrow 0$ (so that $\varepsilon_h \rightarrow 0$, $\Delta w \rightarrow 0$, and $|w_h^* - \bar{w}| \rightarrow 0$) gives $F(\bar{w}) \geq \sup_{\Gamma} F$. Since trivially $F(\bar{w}) \leq \sup_{\Gamma} F$, we conclude $F(\bar{w}) = \sup_{\Gamma} F$, hence $\bar{w} \in \arg \max_{\Gamma} F$. If the maximizer of F on Γ is unique, then the entire sequence w_h^* converges to w^* . \square

We now state the main convergence result in the next theorem.

THEOREM 6.6 (Main convergence of the full scheme). *Fix $x_0 \in \Omega$ and choose $x_0^h \in \Omega^h$ with $x_0^h \rightarrow x_0$ as $h \rightarrow 0$; under (5.4) and Assumption 4.5, the numerical outer-optimized value satisfies*

$$V_h(x_0^h, t_0^-) := \max_{w_h \in \Gamma_h} V_h(x_0^h, w_h, t_0^-) \rightarrow \max_{w \in \Gamma} V(x_0, w, t_0^-) =: V(x_0, t_0^-),$$

and any $w_h^ \in \arg \max_{\Gamma_h} V_h(x_0^h, \cdot, t_0^-)$ has accumulation points in $\arg \max_{\Gamma} V(x_0, \cdot, t_0^-)$ (with $w_h^* \rightarrow w^*$ if the maximizer is unique).*

This follows from Lemmas 6.4–6.5.

7 Generality and extensions Our monotone integration scheme is built from three structural ingredients: (i) translation-invariant inter-decision increments specified via a closed-form CF, (ii) a Fourier-trained transition density $\hat{g} \geq 0$, and (iii) non-negative quadrature/interpolation weights. Thus, for each fixed auxiliary parameter w , the inner Bellman recursion retains the same monotonicity/stability/consistency properties established in Lemmas 6.1–6.4. In particular, for each fixed w , the analysis up to Lemma 6.4 is agnostic to the specific choice of terminal functional $\Phi(\cdot, w)$, which enters only through the terminal condition and the boundary conditions.

Accordingly, our results apply to optimization problems of the form

$$\sup_{w \in \mathcal{W}} \sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} [\Phi(W_T, w)].$$

The outer-maximization convergence argument extends whenever \mathcal{W} is compact (or can be localized to one) and $w \mapsto \Phi(\cdot, w)$ is uniformly continuous on \mathcal{W} : there exists a modulus of continuity ω with $\omega(r) \downarrow 0$ as $r \downarrow 0$ such that

$$\sup_{x \in \Omega} |\Phi(W(x), w_1) - \Phi(W(x), w_2)| \leq \omega(|w_1 - w_2|), \quad \forall w_1, w_2 \in \mathcal{W}.$$

In this case, the outer grid-approximation error is controlled by $\omega(\Delta w)$ (cf. Lemma 6.5).

In many applications, $\Phi(W_T, w)$ admits a natural reward–risk decomposition that highlights a broad class of tail and shortfall criteria encompassed by the same $\sup_w \mathbb{E}[\Phi(\cdot, w)]$ structure. Specifically, consider terminal functionals of the form

$$\Phi(W_T, w) = R(W_T) + \gamma \varphi(Z, w), \quad Z = \zeta(W_T).$$

Here, $R(W_T)$ is a reward term (e.g. $R(W_T) = W_T$), $\gamma > 0$ is a scalarization parameter, and Z is a real-valued variable obtained from W_T via a prescribed measurable mapping $\zeta : \mathbb{R}_+ \rightarrow \mathbb{R}$ (e.g. a loss/shortfall metric derived from the aggregate level W_T). The risk component is encoded through the integrand $\varphi(Z, w)$, which in many cases admits an auxiliary-variable representation of the form $\rho(Z) = \sup_{w \in \mathcal{W}} \mathbb{E}[\varphi(Z, w)]$, so that a terminal reward–risk criterion of the form $\sup_{\mathcal{U}_0 \in \mathcal{A}} \{\mathbb{E}[R(W_T)] + \gamma \rho(Z)\}$ can be written equivalently as $\sup_{w \in \mathcal{W}} \sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} [R(W_T) + \gamma \varphi(Z, w)]$, which matches the $\sup_w \mathbb{E}[\Phi(\cdot, w)]$ structure treated here (with $\Phi(W_T, w) = R(W_T) + \gamma \varphi(Z, w)$).

Beyond CVaR/expected shortfall [27], this class also includes buffered probability of exceedance [22, 7] and optimized certainty equivalents / shortfall-type convex risk measures [4], which can be cast into the same auxiliary-variable $\sup_w \mathbb{E}[\varphi(\cdot, w)]$ representation (possibly after taking a negative when originally posed via an infimum over w). More generally, piecewise-linear (kinked) terminal penalties in Z (and hence in W_T via $Z = \zeta(W_T)$) with breakpoints parameterized by w also fit this setting [4].

Higher dimensions. Finally, the Fourier-trained nonnegative kernel construction, the nonnegative-weight quadrature discretization, and the Barles–Souganidis convergence framework extend to dimensions $d > 2$ under the same structural assumptions, with computational scaling corresponding to d -dimensional quadrature/FFT-based convolutions.

8 Numerical experiments This section reports numerical experiments addressing two separate questions: (i) the approximation accuracy of the Fourier-trained transition kernel \hat{g} when the one-step increment law is specified only through a closed-form CF; and (ii) the performance of the resulting monotone 2D integration scheme when \hat{g} is used to solve a representative multi-period mean-CVaR optimization problem. Subsection 8.1 addresses (i) using a fully coupled 2D jump–diffusion test law with synthetic parameters. Subsection 8.2 then presents a Defined Contribution (DC)

mean-CVaR portfolio illustration for working years (accumulation), where the same increment law is calibrated to long-horizon market data.

Error metrics. Let $f_1, f_2 \in L_p(\mathbb{R}^2)$, $p \in \{1, 2\}$. We report the truncated L_p error $L_p(f_1, f_2) = \int_{[-A, A]^2} |f_1(x) - f_2(x)|^p dx$, for $A > 0$ chosen sufficiently large so that truncation has negligible effect at the displayed precision. We also report the maximum pointwise error $\text{MPE}(f_1, f_2) = \max_{1 \leq k \leq K} |f_1(x_k) - f_2(x_k)|$ over a fixed set of evaluation points $\{x_k\}_{k=1}^K$. Among these, the L_2 error is the primary metric, consistent with the L_2 analysis in Section 4.

8.1 Kernel learning accuracy

8.1.1 2D Kou jump–diffusion We take $G(\eta)$ to be the one-step CF of a fully coupled 2D Kou jump–diffusion model, written in the Lévy-exponent form $G(\eta) = \exp(\Delta t \Psi(\eta))$ [15, 18]. The characteristic exponent $\Psi(\eta_s, \eta_b)$ is given by

$$(8.1) \quad \begin{aligned} \Psi(\eta_s, \eta_b) = & i(\eta_s \mu_s^* + \eta_b \mu_b^*) - \frac{1}{2}(\sigma_s^2 \eta_s^2 + \sigma_b^2 \eta_b^2 + 2\rho \sigma_s \sigma_b \eta_s \eta_b) \\ & + \lambda^s(\varphi_s(\eta_s) - 1) + \lambda^b(\varphi_b(\eta_b) - 1) + \lambda^c(\varphi_c(\eta_s, \eta_b) - 1), \quad |\rho| < 1, \end{aligned}$$

with idiosyncratic jump CFs

$$\varphi_s(\eta_s) = p_s \frac{\eta_1^s}{\eta_1^s - i\eta_s} + (1 - p_s) \frac{\eta_2^s}{\eta_2^s + i\eta_s}, \quad \varphi_b(\eta_b) = p_b \frac{\eta_1^b}{\eta_1^b - i\eta_b} + (1 - p_b) \frac{\eta_2^b}{\eta_2^b + i\eta_b},$$

and a common (co-)jump component

$$\varphi_c(\eta_s, \eta_b) = \varphi_{c,s}(\eta_s) \varphi_{c,b}(\eta_b), \quad \begin{cases} \varphi_{c,s}(\eta_s) = p_{c,s} \frac{\eta_{1,c,s}}{\eta_{1,c,s} - i\eta_s} + (1 - p_{c,s}) \frac{\eta_{2,c,s}}{\eta_{2,c,s} + i\eta_s}, \\ \varphi_{c,b}(\eta_b) = p_{c,b} \frac{\eta_{1,c,b}}{\eta_{1,c,b} - i\eta_b} + (1 - p_{c,b}) \frac{\eta_{2,c,b}}{\eta_{2,c,b} + i\eta_b}. \end{cases}$$

The compensated drifts are

$$\mu_s^* = \mu^s - \frac{(\sigma^s)^2}{2} - \lambda^s \kappa^s - \lambda^c \kappa_s^c, \quad \mu_b^* = \mu^b - \frac{(\sigma^b)^2}{2} - \lambda^b \kappa^b - \lambda^c \kappa_b^c,$$

with compensators

$$\kappa^z = \varphi_z(-i) - 1, \quad z \in \{s, b\}, \quad \kappa_s^c = \varphi_c(-i, 0) - 1, \quad \kappa_b^c = \varphi_c(0, -i) - 1.$$

The synthetic parameters are chosen as follows: $\mu^s = 0.08$, $\sigma^s = 0.03$, $\lambda^s = 0.6$, $p_{\text{up}}^s = 0.4$, $\eta_1^s = 6.5$, $\eta_2^s = 6.5$, $\rho_{sb} = 0.05$; $\mu^b = 0.02$, $\sigma^b = 0.04$, $\lambda^b = 0.8$, $p_{\text{up}}^b = 0.5$, $\eta_1^b = 20.5$, $\eta_2^b = 22.5$; $\lambda^c = 0.2$, $p_{\text{up},s}^c = 0.5$, $\eta_{1,s}^c = 25$, $\eta_{2,s}^c = 30$, $p_{\text{up},b}^c = 0.6$, $\eta_{1,b}^c = 20$, $\eta_{2,b}^c = 35$. For concreteness, we set $\Delta t = 1$ in the training experiments.

8.1.2 Training setup and results We follow the Fourier-domain sampling and truncation prescriptions in Sections 4.2–4.3. We choose η' so that the tail conditions in (4.11) hold with tolerance $\varepsilon = 10^{-6}$; in all experiments we use a conservative value $\eta' = 80$ per axis (increasing η' further does not change the reported results at the shown precision). Table 8.1 summarizes the training hyperparameters.

N	P	# epochs ₁	# epochs ₂	ℓ_1	ℓ_2	batchsize
60	10^6	20	100	0.04	0.00025	1024

Table 8.1: Hyperparameters for Fourier-domain neural-network training.

We use a two-stage optimizer schedule as in [11]: an initial exploration phase (AMSGrad [31]) followed by a refinement phase (Adam [16]). All runs were conducted in Python/TensorFlow; hardware details do not materially affect the reported

accuracy metrics. In this example, both Re, G and Im, G oscillate in the frequency variable η . These oscillations may be rapid over moderate intervals (creating many local minima and noisy gradients) or relatively smooth but spread over a wide domain (reducing sampling and training efficiency). To improve conditioning, we apply a simple affine rescaling of the CF targets as in Remark 5.1 of [11]; the learned kernel is then mapped back and reported in the original coordinates.

Figure 8.1 shows representative slices of G and \widehat{G} , together with a corresponding slice of \widehat{g} . Errors are reported in Table 8.2, where the principal L_2 metric is highlighted. In this illustrative run, the fitted CF yields L_2 errors on the order of 10^{-7} ¹.

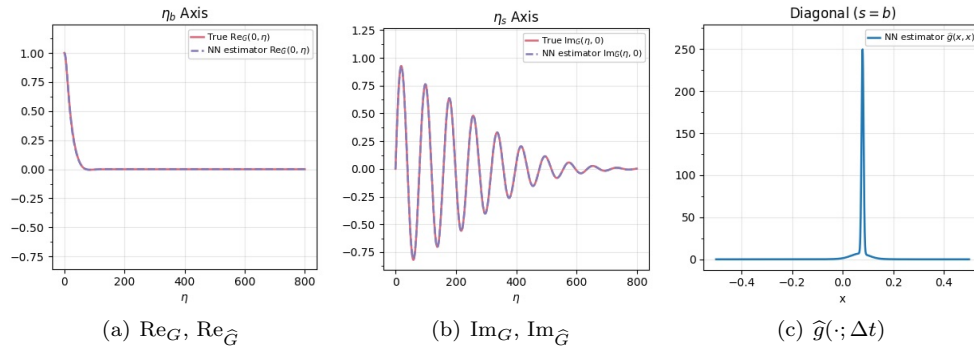


Fig. 8.1: 2D Kou jump–diffusion test case. Panels (a)–(b) show representative 1D slices of the target CF G and the fitted CF \widehat{G} , and panel (c) shows a representative slice of the trained density \widehat{g} . To improve conditioning for oscillatory G , the CF targets in the loss function are affinely rescaled as in Remark 5.1 of [11]; the plotted slices are shown in the original (unrescaled) coordinates.

$L_1(\text{Re}_G, \text{Re}_{\widehat{G}})$	$L_2(\text{Re}_G, \text{Re}_{\widehat{G}})$	$\text{MPE}(\text{Re}_G, \text{Re}_{\widehat{G}})$	$L_1(\text{Im}_G, \text{Im}_{\widehat{G}})$	$L_2(\text{Im}_G, \text{Im}_{\widehat{G}})$	$\text{MPE}(\text{Im}_G, \text{Im}_{\widehat{G}})$
1.1369e-05	9.2574e-07	3.6886e-04	1.0682e-05	7.4543e-08	2.8006e-04

Table 8.2: CF estimation errors for the 2D Kou jump–diffusion test case.

8.2 Mean-CVaR optimization (DC plan) We now specialize the generic 2-D mean-CVaR control problem to a DC accumulation illustration. At each intervention time, as is common in accumulation plans (see, e.g. [7, 14]), the aggregate account value is allocated between the two assets with no borrowing and no negative positions, so the balance invested in each asset remains nonnegative. In this setting, the state $X_t = (S_t, B_t)$ represents the log-balances invested in an equity index (component s) and a bond index (component b), respectively. The corresponding aggregate account value is $W_t = e^{S_t} + e^{B_t}$, so the terminal outcome W_T in the mean-CVaR objective is the total retirement balance at horizon T .

Between decision times, the increment $\Delta X = ((\Delta S)_m, (\Delta B)_m)$ represents the joint log-return over one period. At each intervention time t_m , the account receives a deterministic (salary) contribution q_m , and the control $u_m \in [0, 1]$ reallocates the post-contribution wealth across the two components as in (3.1)–(3.3).

¹We also tested bivariate Variance-Gamma, bivariate NIG, and 2D Merton specifications; the qualitative kernel-fit accuracy and control outcomes are similar and are omitted for brevity.

8.2.1 Calibration summary To provide a concrete mean-CVaR control illustration with empirically realistic dynamics, we calibrate the 2D Kou model of Subsection 8.1 to long-horizon U.S. market data. Specifically, we use monthly total return series from the Center for Research in Security Prices (CRSP) over 1926:01–2024:12². In the DC illustration below, the two components correspond to equity and bond log-returns, respectively. The resulting annualized parameter estimates are reported in Table 8.3. Dependence is introduced via both the diffusion correlation and the common-jump component.

Method	μ^s	σ^s	λ^s	p_{up}^s	η_1^s	η_2^s	ρ_{sb}
Component s (Idiosyncratic; equity index)							
	0.0898	0.1326	0.5960	0.373	6.701	6.634	(see below)
Method	μ^b	σ^b	λ^b	p_{up}^b	η_1^b	η_2^b	ρ_{sb}
Component b (Idiosyncratic; bond index)							
	0.0204	0.0466	0.9495	0.479	20.764	22.551	0.0721
Common Jump							
Method	λ^c	$p_{\text{up},s}^c$	$\eta_{1,s}^c$	$\eta_{2,s}^c$	$p_{\text{up},b}^c$	$\eta_{1,b}^c$	$\eta_{2,b}^c$
	0.1010	0.300	9.825	7.146	0.500	16.982	19.914

Table 8.3: Estimated annualized parameters for the 2D Kou model with idiosyncratic and common jumps, calibrated from CRSP equity index and 10-year Treasury total returns; monthly, 1926:01–2024:12. Threshold technique from [9].

8.2.2 Accumulation scenario We now illustrate the full pipeline (kernel-learning + strictly monotone 2D integration) in a multi-period mean-CVaR control problem with discrete interventions, using a defined DC retirement-accumulation setting. This DC example is used only as an application-level illustration (with interpretable interventions and controls); the kernel-learning construction and convergence analysis are not specific to finance.

To illustrate the accumulation phase of a DC plan, we consider a 35-year-old investor with an annual salary of \$100,000. The total contribution to the plan account is 20% of salary each year. The investor plans to retire at age 65, yielding a 30-year savings horizon [14]. The scenario is summarized in Table 8.4. In this illustration we set $\Delta t = 1$ year and use the calibrated parameters in Table 8.3. We train \hat{g} using the same setup as in Subsection 8.1.2; the resulting CF-fit errors are of comparable magnitude to Table 8.2, so we omit repeated slices/metrics for brevity.

Investment horizon (years)	30
Initial investment W_0	0
Cash contributions	\$20,000/year
Rebalancing frequency	yearly

Table 8.4: DC accumulation illustration. Cash contributions are made at $t_m = 0, 1, \dots, 29$ years. Model parameters are given in Table 8.3.

²The results presented here were calculated based on data from Historical Indexes, © 2015 Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business.

8.2.3 Convergence results and efficient frontier We choose a computational domain large enough that boundary truncation effects are negligible [33, 8]. For the calibrated parameters (Table 8.3) and the DC accumulation scenario (Table 8.4), the baseline domains are listed in Table 8.5. The padded bounds $(s_{\min}^\dagger, s_{\max}^\dagger)$ and $(b_{\min}^\dagger, b_{\max}^\dagger)$ are set according to (5.1). Unless otherwise stated, refinement levels and discretization parameters are given in Table 8.6.

s_{\min}	s_{\max}
$\ln(10^9) - 8$	$\ln(10^9) + 8$
b_{\min}	b_{\max}
$\ln(10^9) - 8$	$\ln(10^9) + 8$
w_{\max}	
1×10^8	

Table 8.5: Computational domains of numerical experiments.

Refine. level	N_s	N_b	N_u	N_w
0	512	512	256	512
1	1024	1024	512	1024
2	2048	2048	1024	2048

Table 8.6: Discretization parameters and grid refinement levels.

Table 8.7 reports a convergence check for the mean-CVaR optimization problem with $\gamma = 10$ and $\alpha = 0.05$. The differences between the two finest refinement levels suggest that the computed objective is accurate to well within 1%. The optimal controls are computed and stored, and then used as inputs to Monte Carlo simulation (2.56×10^6 simulations) for validation (reported in the same table).

Refine. level	Our method			Monte Carlo		
	$\mathbb{E}[W_T]$	CVaR (5%)	w^*	$\mathbb{E}[W_T]$	CVaR (5%)	Median $[W_T]$
0	2765.81	624.08	733.86	2760.64 (6.1)	624.50	1481.44
1	2768.71	624.37	734.12	2768.12 (6.1)	624.37	1481.34
2	2769.90	624.42	734.25	2770.18 (6.1)	624.41	1481.23

Table 8.7: Convergence test, mean-CVaR with $\gamma = 10$ and $\alpha = 0.05$. Parameters in Table 8.3. Brackets show half-widths of 99% confidence intervals. Units: thousands of currency units.

Figure 8.2 shows the optimal allocation weight in component s as a function of time and realized wealth. Early in the horizon the policy is near full allocation to component s ; as time advances, allocations become more state-dependent, with a clear de-risking region appearing around the wealth range associated with the optimal threshold $w^* \approx 0.73$ million (Table 8.7). For wealth well below this level, the policy increases exposure to component s to improve tail outcomes; for wealth near the threshold, the policy becomes more conservative; and for sufficiently high wealth, the CVaR term is less binding and the allocation again tilts toward higher expected growth.

To construct the efficient frontier, we vary the scalarization parameter $\gamma \in [0.01, 1000]$ and repeat the optimization at each γ (all results on the finest refinement level). Figure 8.3 plots the resulting frontier in the $(\text{CVaR}_\alpha(W_T), \mathbb{E}[W_T])$ plane with $\alpha = 0.05$. Since CVaR is applied to terminal wealth in this illustration, larger $\text{CVaR}_\alpha(W_T)$ corresponds to better downside performance. Increasing γ shifts emphasis toward improving tail outcomes (higher CVaR) at the expense of expected terminal wealth.

8.2.4 Robustness checks

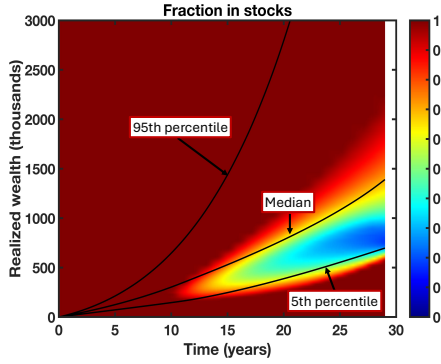


Fig. 8.2: Pre-commitment mean-CVaR optimal control heat map (allocation weight in component s ; interpreted as equity in the DC illustration).

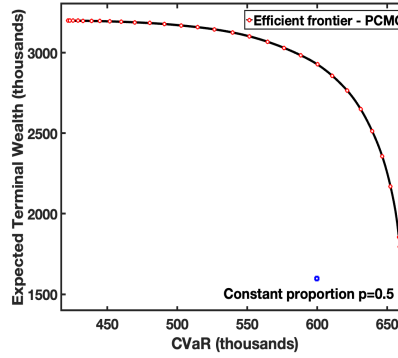


Fig. 8.3: Efficient frontier of mean-CVaR with $\alpha = 0.05$, computed on the finest refinement level (DC illustration).

Impact of spatial domain sizes To validate that the baseline domains in Table 8.5 are sufficiently large, we repeat the experiment reported in Table 8.7 on (i) a moderately larger domain and (ii) a smaller domain, adjusting N_s, N_b to keep $\Delta s, \Delta b$ unchanged. Tables 8.8–8.9 show that enlarging the domain has negligible impact at the reported precision, while shrinking the domain introduces small discrepancies (as expected from truncation effects).

Level	Larger domain			Table 8.7		
	$\mathbb{E}[W_T]$	CVaR (5%)	w^*	$\mathbb{E}[W_T]$	CVaR (5%)	w^*
0	2765.80	624.08	733.86	2765.81	624.08	733.86
1	2768.71	624.37	734.12	2768.71	624.37	734.12
2	2769.90	624.42	734.25	2769.90	624.42	734.25

Table 8.8: Results using a larger spatial domain: $z_{\min} = \ln(10^5) - 10$, $z_{\max} = \ln(10^5) + 10$ (baseline domain $z_{\min} = \ln(10^5) - 8$, $z_{\max} = \ln(10^5) + 8$, Table 8.5) for $z \in \{s, b\}$.

Level	Smaller domain			Table 8.7		
	$\mathbb{E}[W_T]$	CVaR (5%)	w^*	$\mathbb{E}[W_T]$	CVaR (5%)	w^*
0	2765.01	624.10	733.86	2765.81	624.08	733.86
1	2767.96	624.39	734.12	2768.71	624.37	734.12
2	2769.19	624.45	734.25	2769.90	624.42	734.25

Table 8.9: Results using a smaller spatial domain: $z_{\min} = \ln(10^5) - 6.25$, $z_{\max} = \ln(10^5) + 6.25$ for $z \in \{s, b\}$; (baseline domain from Table 8.5).

Impact of boundary conditions We repeat the experiment in Table 8.7 using the constant boundary conditions proposed in [34, 10]. Table 8.10 shows that, on the baseline domain, the constant boundary results are close to, but not fully identical with, those obtained using our asymptotic boundary conditions (3.17). As a robustness check, Table 8.11 repeats the constant boundary experiment on a larger domain,

which reduces the discrepancy, consistent with constant boundary conditions requiring larger domains to reach comparable accuracy.

Level	Constant boundary			Table 8.7		
	$\mathbb{E}[W_T]$	CVaR (5%)	w^*	$\mathbb{E}[W_T]$	CVaR (5%)	w^*
0	2765.77	624.07	733.86	2765.81	624.08	733.86
1	2768.64	624.38	734.12	2768.71	624.37	734.12
2	2769.76	624.42	734.25	2769.90	624.42	734.25

Table 8.10: Constant boundary conditions on the baseline domain, compared with our asymptotic boundary conditions (3.17).

Level	Constant boundary			Tab. 8.7		
	$\mathbb{E}[W_T]$	CVaR (5%)	W^*	$\mathbb{E}[W_T]$	CVaR (5%)	W^*
0	2765.80	624.08	733.86	2765.81	624.08	733.86
1	2768.71	624.37	734.12	2768.71	624.37	734.12
2	2769.90	624.42	734.25	2769.90	624.42	734.25

Table 8.11: Constant boundary conditions on a larger domain, compared with our asymptotic boundary conditions (3.17).

9 Conclusion We developed a strictly monotone 2D integration scheme for multi-period mean-CVaR optimization, a representative class of reward-risk stochastic control problems with discrete interventions, in settings where the between-intervention increment law is specified via a closed-form CF with mild Fourier-tail decay.

The key computational ingredient is a Fourier-trained, nonnegative Gaussian-mixture transition kernel, which enables direct composite-quadrature evaluation of the Bellman convolution and an efficient FFT implementation. Our error analysis is conducted in Fourier space, leveraging the explicit Fourier form of both the Gaussian-mixture kernel and the target CF: we derive Fourier-domain L_2 error estimates and translate them into real-space bounds, which are then used to establish ℓ_∞ stability, consistency, and pointwise convergence as the discretization and kernel-approximation parameters vanish.

Numerical experiments include (i) a fully coupled 2-D jump-diffusion test law to assess kernel-learning accuracy, and (ii) a DC mean-CVaR optimization example representing accumulation over working years, calibrated to long-horizon data. Together, they demonstrate practical accuracy and robustness. More broadly, the Fourier-trained kernel construction and the Fourier-to-real-space convergence framework apply to other reward-risk stochastic control problems with translation-invariant kernels, and can be extended beyond two dimensions, subject to the standard computational considerations. Future work includes time-consistent optimization formulations and adaptive Fourier training strategies that preserve strict monotonicity.

Appendix A. Proof of Lemma 3.2.

For each $w \geq 0$, define $F(w) := \sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} [W_T + \gamma(w + \frac{1}{\alpha} \min(W_T - w, 0))]$.

By Remark 3.1, $K := \sup_{\mathcal{U}_0 \in \mathcal{A}} \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0^-} [W_T] < \infty$, hence $F(w)$ is finite for each $w \geq 0$.

Next, use $\min(W_T - w, 0) \leq W_T - w$ to obtain, for any admissible \mathcal{U}_0 ,

$$W_T + \gamma \left(w + \frac{1}{\alpha} \min(W_T - w, 0) \right) \leq \left(1 + \frac{\gamma}{\alpha} \right) W_T + \gamma \left(1 - \frac{1}{\alpha} \right) w.$$

Taking expectations and then $\sup_{\mathcal{U}_0 \in \mathcal{A}}$ yields the uniform bound

$$F(w) \leq \left(1 + \frac{\gamma}{\alpha} \right) K + \gamma \left(1 - \frac{1}{\alpha} \right) w.$$

Since $1 - \frac{1}{\alpha} < 0$, the right-hand side tends to $-\infty$ as $w \rightarrow \infty$, hence $F(w) \rightarrow -\infty$ as $w \rightarrow \infty$. Finally, for each fixed \mathcal{U}_0 , the map $w \mapsto w + \frac{1}{\alpha} \min(W_T - w, 0)$ is Lipschitz in w (uniformly in W_T), so $w \mapsto \mathbb{E}_{\mathcal{U}_0}^{x_0, t_0} [\cdot]$ is continuous; therefore $F(\cdot) = \sup_{\mathcal{U}_0} (\cdot)$ is Lipschitz and in particular continuous. Thus, there exists $\bar{w} < \infty$ such that $\sup_{w \geq 0} F(w) = \max_{w \in [0, \bar{w}]} F(w)$, and by continuity the maximum is attained at some $w^*(x_0) \in [0, \bar{w}] \subset [0, \infty)$.

Appendix B. Proof of Lemma 4.2. Let $\Delta(\eta) := G(\eta) - \widehat{G}(\eta; \widehat{\theta}^*)$.

(i) *L₂ bound.* By (4.6) applied to $f = g - \widehat{g}$,

$$\int_{\mathbb{R}^2} |g - \widehat{g}|^2 = \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} |\Delta(\eta)|^2 d\eta.$$

Split $\mathbb{R}^2 = D_\eta \cup (\mathbb{R}^2 \setminus D_\eta)$. On $\mathbb{R}^2 \setminus D_\eta$, $|\Delta|^2 \leq 2(|G|^2 + |\widehat{G}|^2)$, so (4.11) yields $\int_{\mathbb{R}^2 \setminus D_\eta} |\Delta|^2 < 4\varepsilon_1$. On D_η , approximate $\int_{D_\eta} |\Delta|^2$ by a composite left-hand quadrature $\sum_{p=1}^P w_p |\Delta(\eta_p)|^2$. Under (4.8), the weights satisfy $0 < w_p \leq C_1/P$ and the quadrature remainder is bounded by $\frac{C' C_1^3}{2P^{1/2}}$ (using (4.12)). Hence

$$\int_{D_\eta} |\Delta|^2 d\eta \leq C_1 \frac{1}{P} \sum_{p=1}^P |\Delta(\eta_p)|^2 + \frac{C' C_1^3}{2P^{1/2}} \leq C_1 \text{Loss}_P(\widehat{\theta}^*) + \frac{C' C_1^3}{2P^{1/2}},$$

and since $\text{Loss}_P(\widehat{\theta}^*) < \varepsilon_2$, this gives (4.13).

(ii) *Pointwise bound.* From the inverse transform (2.3) and $|a + ib| \leq |a| + |b|$,

$$|g(y) - \widehat{g}(y; \widehat{\theta}^*)| \leq \frac{1}{(2\pi)^2} \int_{\mathbb{R}^2} (|\text{Re}_\Delta(\eta)| + |\text{Im}_\Delta(\eta)|) d\eta.$$

Split $\mathbb{R}^2 = D_\eta \cup (\mathbb{R}^2 \setminus D_\eta)$. On the tail,

$$\int_{\mathbb{R}^2 \setminus D_\eta} (|\text{Re}_\Delta| + |\text{Im}_\Delta|) d\eta \leq \int_{\mathbb{R}^2 \setminus D_\eta} (|\text{Re}_G| + |\text{Im}_G|) d\eta + \int_{\mathbb{R}^2 \setminus D_\eta} (|\text{Re}_{\widehat{G}}| + |\text{Im}_{\widehat{G}}|) d\eta \leq 2\varepsilon_1 d\eta.$$

On D_η , applying the same quadrature bound with weights $w_p \leq C_1/P$ gives

$$\int_{D_\eta} (|\text{Re}_\Delta| + |\text{Im}_\Delta|) d\eta \leq \sum_{p=1}^P w_p (|\text{Re}_\Delta(\eta_p)| + |\text{Im}_\Delta(\eta_p)|) + \frac{C' C_1^3}{2P^{1/2}} \leq C_1 R_P(\widehat{\theta}^*) + \frac{C' C_1^3}{2P^{1/2}}.$$

Since $R_P(\widehat{\theta}^*) < \varepsilon_3$, (4.14) follows.

Appendix C. Proof of Lemma 4.3. Let $\lambda := c_0 \Delta t > 0$ and fix $\eta' \geq R$. For any $\eta \in \mathbb{R}^2$ with $\|\eta\|_\infty > \eta'$ we have $|\eta| > \eta'$. Under Assumption 2.1 (A4),

$$|G(\eta; \Delta t)|^p \leq \exp(-p\lambda|\eta|^\alpha), \quad p \in \{1, 2\}.$$

Hence, for $p \in \{1, 2\}$, noting $\mathbb{R}^2 \setminus [-\eta', \eta']^2 = \{\eta \in \mathbb{R}^2 \mid \|\eta\|_\infty > \eta'\}$,

$$\int_{\{\|\eta\|_\infty > \eta'\}} |G(\eta; \Delta t)|^p d\eta \leq \int_{\{|\eta| > \eta'\}} e^{-p\lambda|\eta|^\alpha} d\eta = 2\pi \int_{\eta'}^\infty r e^{-p\lambda r^\alpha} dr.$$

Make the substitution $u = p\lambda r^\alpha$; then $r = (u/(p\lambda))^{1/\alpha}$ and $dr = \frac{1}{\alpha}(p\lambda)^{-1/\alpha} u^{\frac{1}{\alpha}-1} du$,

$$2\pi \int_{\eta'}^{\infty} r e^{-p\lambda r^\alpha} dr = \frac{2\pi}{\alpha} (p\lambda)^{-2/\alpha} \Gamma\left(\frac{2}{\alpha}, p\lambda \eta'^\alpha\right),$$

where $\Gamma(a, z)$ is the upper incomplete gamma function. It is standard that for any $a > 0$ and any $\theta \in (0, 1)$ there exists $C(a, \theta) > 0$ such that $\Gamma(a, z) \leq C(a, \theta) e^{-\theta z}$ for all $z \geq 0$ [1]. Taking $\theta = \frac{1}{2}$ and $z = p\lambda \eta'^\alpha$ in this bound, noting $\lambda = c_0 \Delta t$, yields

$$\int_{\{\|\eta\|_\infty > \eta'\}} |G(\eta; \Delta t)|^p d\eta \leq C_G^{(p)} \exp(-c_G^{(p)} \eta'^\alpha),$$

for constants $C_G^{(p)} = C(a, \theta) \frac{2\pi}{\alpha} (p c_0 \Delta t)^{-2/\alpha}$ and $c_G^{(p)} := \frac{1}{2} p c_0 \Delta t > 0$, independent of $\eta' \geq R$. This concludes the Fourier domain truncation error bound related to G .

Regarding \widehat{G} , for $\eta = (\eta_s, \eta_b) \in \mathbb{R}^2$, set $r := \sigma_n^s \eta_s$ and $q := \sigma_n^b \eta_b$. We then have

$$\eta^\top \Sigma_n \eta = r^2 + q^2 + 2\rho_n r q \geq r^2 + q^2 - 2|\rho_n| |r q| \geq (1 - |\rho_n|) (r^2 + q^2) \geq (1 - \bar{\rho}) \sigma_{\min}^2 |\eta|^2.$$

Hence, $|\widehat{G}(\eta; \theta)| \leq \sum_{n=1}^N |\beta_n| e^{-\frac{1}{2} \eta^\top \Sigma_n \eta} \leq N \bar{\beta} \exp\left(-\frac{(1-\bar{\rho})\sigma_{\min}^2}{2} |\eta|^2\right)$.

Let $\sigma^2 := (1 - \bar{\rho}) \sigma_{\min}^2$. For $p \in \{1, 2\}$,

$$\int_{\{\|\eta\|_\infty > \eta'\}} |\widehat{G}(\eta; \theta)|^p d\eta \leq (N \bar{\beta})^p \int_{\{|\eta| > \eta'\}} e^{-\frac{p\sigma^2}{2} |\eta|^2} d\eta = (N \bar{\beta})^p \frac{2\pi}{p\sigma^2} e^{-\frac{p\sigma^2}{2} \eta'^2}.$$

Choosing $c_{\widehat{G}}^{(p)} := \frac{p\sigma^2}{2} = \frac{p(1-\bar{\rho})\sigma_{\min}^2}{2}$ and $C_{\widehat{G}}^{(p)} := (N \bar{\beta})^p \frac{2\pi}{p\sigma^2}$ concludes the proof.

Appendix D. Proof of Corollary 4.6. Choose $\eta'(h)$ so that the tail terms in Lemma 4.2 satisfy $\varepsilon_1 \leq c_1 h^{1+\kappa}$ (Assumption 4.5). Assume $\varepsilon_2 \leq c_2 h^{1+\kappa}$, $\varepsilon_3 \leq c_3 h^{1+\kappa}$, and $P \geq c_P h^{-2(1+\kappa)}$, so $P^{-1/2} \leq c_P^{-1/2} h^{1+\kappa}$. Applying Lemma 4.2 (i) gives $\|g - \widehat{g}\|_{L_2}^2 \leq C h^{1+\kappa}$, hence $\|g - \widehat{g}\|_{L_2} \leq C h^{(1+\kappa)/2}$. Applying Lemma 4.2 (ii) gives (4.15).

Appendix E. Proof of Lemma 6.1. By (6.3), there exists $h_0 > 0$ such that for all $0 < h \leq h_0$,

$$(E.1) \quad \Delta s \Delta b \sum_{l,d} \varphi_{l,d} |\widehat{g}_{l-k, d-j}| \leq 1 + \varepsilon(h) \leq e^{\varepsilon(h)}, \quad \text{for all } (s_k, b_j) \in \Omega_{\text{in}},$$

with $\varepsilon(h) = C_\varepsilon h^{1+\kappa}$ and C_ε independent of h . Define

$$\kappa_{\max} := \frac{1}{\Delta t} \max \{ \ln |G(-i, 0; \Delta t)|, \ln |G(0, -i; \Delta t)|, \ln |G(-i, -i; \Delta t)|, 0 \}.$$

We claim that, for $m = 0, \dots, M$,

$$(E.2) \quad \|V_c^{m,-}\|_\infty \leq e^{(M-m)(\Delta t \kappa_{\max} + \varepsilon(h))} \|V_c^{M,-}\|_\infty.$$

The base case $m = M$ is immediate.

Assume (E.2) holds at $m+1$. For boundary nodes $(s_k, b_j) \in \Omega_{\text{out}}$, (5.7) and (3.16) give

$$|V_{k,j,c}^{m,-}| = |V_{k,j,c}^{m,+}| = |G(-i a(s_k, b_j); \Delta t)| |V_{k,j,c}^{m+1,-}| \leq e^{\Delta t \kappa_{\max}} \|V_c^{m+1,-}\|_\infty.$$

For interior nodes $(s_k, b_j) \in \Omega_{\text{in}}$, (5.8) and (E.1) yield

$$|V_{k,j,c}^{m,+}| \leq \left(\Delta s \Delta b \sum_{l,d} \varphi_{l,d} |\widehat{g}_{l-k, d-j}| \right) \|V_c^{m+1,-}\|_\infty \leq e^{\varepsilon(h)} \|V_c^{m+1,-}\|_\infty.$$

The intervention step (5.9) uses bilinear interpolation with nonnegative weights summing to one, followed by a max over controls, hence it is non-expansive in ℓ_∞ : $\|V_c^{m,-}\|_\infty \leq \|V_c^{m,+}\|_\infty$. Therefore

$$\|V_c^{m,-}\|_\infty \leq e^{\Delta t \kappa_{\max} + \varepsilon(h)} \|V_c^{m+1,-}\|_\infty,$$

which, combined with the induction hypothesis, implies (E.2).

Next,

$$(E.3) \quad \|V_c^{M,-}\|_\infty \leq \sup_{x \in \Omega, w \in \Gamma} |V(x, w, T^-)| < \infty,$$

since $\Omega \times \Gamma$ is compact and $V(\cdot, \cdot, T^-)$ from (3.8) is continuous. Choose h_0 smaller if needed so that $M \varepsilon(h) \leq 1$ for $0 < h \leq h_0$. Then for all m ,

$$(E.4) \quad e^{(M-m)(\Delta t \kappa_{\max} + \varepsilon(h))} \leq e^{T \kappa_{\max} + 1}.$$

Combining (E.2)–(E.4) yields the stated uniform ℓ_∞ bound.

Appendix F. Proof of Lemma 6.2.

Proof. The equality $\mathcal{D}_h = \mathcal{D}$ is immediate at $m = M$ (terminal payoff) and on Ω_{out}^h for $m = M - 1, \dots, 0$. We therefore focus on the interior.

Fix $(s_k, b_j) \in \Omega_{\text{in}}^h$ and write $x_{k,j} = (s_k, b_j)$. Let $\phi_{m+1}(y) := \phi(y, w_c, t_{m+1}^-)$ and define the continuous propagator

$$\mathcal{Q}(x) := \mathbf{1}_{\Omega_{\text{in}}}(x) \int_{\Omega} \phi_{m+1}(y) g(y - x; \Delta t) dy + \mathbf{1}_{\Omega_{\text{out}}}(x) G(-i a(x); \Delta t) \phi_{m+1}(x),$$

$$\tilde{\mathcal{Q}}(x) := \mathbf{1}_{\Omega_{\text{in}}}(x) \int_{\Omega} \phi_{m+1}(y) \hat{g}(y - x; \hat{\theta}^*) dy + \mathbf{1}_{\Omega_{\text{out}}}(x) G(-i a(x); \Delta t) \phi_{m+1}(x).$$

Define the discrete propagation values (using the same \hat{g} on Ω_{in}^h and the boundary update on Ω_{out}^h)

$$\begin{aligned} \mathcal{Q}_h(x_{l',d'}) &:= \mathbf{1}_{\Omega_{\text{in}}^h}(x_{l',d'}) \Delta s \Delta b \sum_{l,d} \varphi_{l,d} \phi(\hat{x}_{l,d,c}^{m+1,-}) \hat{g}(y_{l,d} - x_{l',d'}; \hat{\theta}^*) \\ &\quad + \mathbf{1}_{\Omega_{\text{out}}^h}(x_{l',d'}) G(-i a(x_{l',d'}); \Delta t) \phi(\hat{x}_{l',d',c}^{m+1,-}), \quad x_{l',d'} \in \Omega^h, \end{aligned}$$

and let $\tilde{\mathcal{Q}}_h(x_{l',d'}) := \tilde{\mathcal{Q}}(x_{l',d'})$ denote nodal samples of $\tilde{\mathcal{Q}}$ on Ω^h .

With $x_{k,j}^{m,+}(u) := x^+(x_{k,j}, q_m, u) \equiv (s_{k,j}^{m,+}(u), b_{k,j}^{m,+}(u))$, $\mathcal{D}(\cdot)$ and $\mathcal{D}_h(\cdot)$ can be written as

$$\mathcal{D}(\hat{x}_{k,j,c}^{m,-}, \phi^{m+1,-}) = \sup_{u \in \mathcal{Z}} \mathcal{Q}(x_{k,j}^{m,+}(u)),$$

$$\mathcal{D}_h(\hat{x}_{k,j,c}^{m,-}, \{\phi(\hat{x}_{l,d,c}^{m+1,-})\}_{l,d}) = \max_{u \in \mathcal{Z}_h} \mathcal{I}[\{\mathcal{Q}_h(x_{l',d'})\}_{l',d'}](x_{k,j}^{m,+}(u)).$$

For any $u \in \mathcal{Z}_h$, set $\xi := x_{k,j}^{m,+}(u)$. Using bilinear interpolation,

$$(F.1) \quad |\mathcal{I}[\mathcal{Q}_h](\xi) - \mathcal{Q}(\xi)| \leq |\mathcal{I}[\mathcal{Q}_h](\xi) - \mathcal{I}[\tilde{\mathcal{Q}}_h](\xi)| + |\mathcal{I}[\tilde{\mathcal{Q}}_h](\xi) - \tilde{\mathcal{Q}}(\xi)| + |\tilde{\mathcal{Q}}(\xi) - \mathcal{Q}(\xi)|.$$

Quadrature term. Since $y \mapsto \phi_{m+1}(y) \hat{g}(y - x)$ is smooth on Ω with uniformly bounded second derivatives (Gaussian mixture kernel; ϕ smooth), the 2-D composite trapezoid rule gives, for each $x_{l',d'} \in \Omega_{\text{in}}^h$,

$$|\mathcal{Q}_h(x_{l',d'}) - \tilde{\mathcal{Q}}_h(x_{l',d'})| \leq C_q h^2,$$

and for $x_{l',d'} \in \Omega_{\text{out}}^h$ the difference is 0 by construction. Hence

$$\|\mathcal{Q}_h - \tilde{\mathcal{Q}}_h\|_{L_\infty(\Omega)} \leq C_q h^2, \quad \Rightarrow \quad |\mathcal{I}[\mathcal{Q}_h](\xi) - \mathcal{I}[\tilde{\mathcal{Q}}_h](\xi)| \leq C_q h^2.$$

Interpolation term. On each region where its definition is smooth (in particular on Ω_{in} and on each boundary subdomain where $a(\cdot)$ is constant), $\tilde{\mathcal{Q}}$ has bounded second derivatives, hence bilinear interpolation yields

$$(F.2) \quad |\mathcal{I}[\tilde{\mathcal{Q}}_h](\xi) - \tilde{\mathcal{Q}}(\xi)| \leq C_{\text{int}} h^2,$$

uniformly in $\xi \in \Omega$, with bounded constant $C_{\text{int}} > 0$ independent of h .

Kernel mismatch term. If $\xi \in \Omega_{\text{out}}$, then $\tilde{\mathcal{Q}}(\xi) = \mathcal{Q}(\xi)$ by definition. If $\xi \in \Omega_{\text{in}}$, using the global pointwise estimate $\|\hat{g} - g\|_{L_\infty(\mathbb{R}^2)} = O(h^{1+\kappa})$ (Assumption 4.5), we have $|\tilde{\mathcal{Q}}(\xi) - \mathcal{Q}(\xi)| = \dots$

$$(F.3) \quad \dots = \left| \int_{\Omega} \phi_{m+1}(y) (\hat{g} - g)(y - \xi) dy \right| \leq \|\phi_{m+1}\|_{L_\infty(\Omega)} |\Omega| \|\hat{g} - g\|_{L_\infty(\mathbb{R}^2)} = C_g h^{1+\kappa}.$$

Combining (F.1)–(F.3) and taking the max over $u \in \mathcal{Z}_h$ gives

$$(F.4) \quad \left| \max_{u \in \mathcal{Z}_h} \mathcal{I}[\mathcal{Q}_h](x_{k,j}^{m,+}(u)) - \max_{u \in \mathcal{Z}_h} \mathcal{Q}(x_{k,j}^{m,+}(u)) \right| \leq C_q h^2 + C_{\text{int}} h^2 + C_g h^{1+\kappa}.$$

Finally, since $u \mapsto x_{k,j}^{m,+}(u)$ is Lipschitz on $[0, 1]$ and $x \mapsto \mathcal{Q}(x)$ is Lipschitz on Ω , the control discretization satisfies

$$(F.5) \quad \left| \sup_{u \in \mathcal{Z}} \mathcal{Q}(x_{k,j}^{m,+}(u)) - \max_{u \in \mathcal{Z}_h} \mathcal{Q}(x_{k,j}^{m,+}(u)) \right| \leq C_u \Delta u = \mathcal{O}(h).$$

Adding (F.4) and (F.5) yields the stated $\mathcal{O}(h + h^{1+\kappa} + h^2)$ bound. \square

REFERENCES

- [1] M. ABRAMOWITZ AND I. A. STEGUN, *Handbook of mathematical functions*, Dover, New York, 1972.
- [2] G. BARLES AND P. SOUGANIDIS, *Convergence of approximation schemes for fully nonlinear equations*, *Asymptotic Analysis*, 4 (1991), pp. 271–283.
- [3] O. E. BARNDORFF-NIELSEN, *Normal inverse gaussian distributions and stochastic volatility modelling*, *Scandinavian Journal of statistics*, 24 (1997), pp. 1–13.
- [4] A. BEN-TAL AND M. TEBoulLE, *An old-new concept of convex risk measures: The optimized certainty equivalent*, *Mathematical Finance*, 17 (2007), pp. 449–476.
- [5] P. CARR, H. GEMAN, D. B. MADAN, AND M. YOR, *The fine structure of asset returns: An empirical investigation*, *The Journal of Business*, 75 (2002), pp. 305–332.
- [6] Z. CHEN AND P. A. FORSYTH, *A numerical scheme for the impulse control formulation for pricing variable annuities with a Guaranteed Minimum Withdrawal Benefit (GMWB)*, *Numerische Mathematik*, 109 (2008), pp. 535–569.
- [7] D.-M. DANG AND C. CHEN, *Multi-period mean-buffered probability of exceedance in Defined Contribution portfolio optimization*, *SIAM Journal on Financial Mathematics*, (2026). to appear.
- [8] D.-M. DANG AND P. FORSYTH, *Continuous time mean-variance optimal portfolio allocation under jump diffusion: A numerical impulse control approach*, *Numerical Methods for Partial Differential Equations*, 30 (2014), pp. 664–698.
- [9] D.-M. DANG AND P. A. FORSYTH, *Better than pre-commitment mean-variance portfolio allocation strategies: A semi-self-financing Hamilton–Jacobi–Bellman equation approach*, *European Journal of Operational Research*, 250 (2016), pp. 827–841.
- [10] D.-M. DANG AND H. ZHOU, *A monotone piecewise constant control integration approach for the two-factor uncertain volatility model*, *IMA Journal of Numerical Analysis*, (2025), <https://doi.org/10.1093/imanum/draf095>.
- [11] R. DU AND D.-M. DANG, *Fourier neural network approximation of transition densities in finance*, *SIAM Journal on Scientific Computing*, 47 (2025), pp. C529–C557.
- [12] F. FANG AND C. OOSTERLEE, *A novel pricing method for European options based on Fourier-Cosine series expansions*, *SIAM Journal on Scientific Computing*, 31 (2008), pp. 826–848.
- [13] C. FILIPPI, G. GUASTAROLA, AND M. G. SPERANZA, *Conditional value-at-risk beyond finance: a survey*, *International Transactions in Operational Research*, 27 (2020), pp. 1277–1319.

- [14] P. A. FORSYTH, *Multiperiod mean Conditional Value at Risk asset allocation: Is it advantageous to be time consistent?*, SIAM Journal on Financial Mathematics, 11 (2020), pp. 358–384.
- [15] S. KEN-ITI, *Lévy processes and infinitely divisible distributions*, vol. 68, Cambridge University Press, 1999.
- [16] D. KINGA, J. B. ADAM, ET AL., *A method for stochastic optimization*, in International conference on learning representations (ICLR), vol. 5, California, 2015.
- [17] P. R. KLEINDORFER AND G. H. SAAD, *Managing disruption risks in supply chains*, Production and Operations Management, 14 (2005), pp. 53–68.
- [18] S. KOU, *A jump diffusion model for option pricing*, Management Science, 48 (2002), pp. 1086–1101.
- [19] J. W. LABADIE, *Optimal operation of multireservoir systems: State-of-the-art review*, Journal of Water Resources Planning and Management, 130 (2004), pp. 93–111.
- [20] Y. LU AND D.-M. DANG, *A semi-Lagrangian ε -monotone Fourier method for continuous withdrawal GMWBs under jump-diffusion with stochastic interest rate*, Numerical Methods for Partial Differential Equations, 40 (2024), p. e23075.
- [21] K. MA AND P. FORSYTH, *An unconditionally monotone numerical scheme for the two-factor uncertain volatility model*, IMA Journal of Numerical Analysis, 37(2) (2017), pp. 905–944.
- [22] A. MAFUSALOV AND S. URYASEV, *Buffered probability of exceedance: Mathematical properties and optimization*, SIAM Journal on Optimization, 28 (2018), pp. 1077–1103.
- [23] G. J. MCLACHLAN, S. X. LEE, AND S. I. RATHNAYAKE, *Finite mixture models*, Annual review of statistics and its application, 6 (2019), pp. 355–378.
- [24] R. C. MERTON, *Option pricing when underlying stock returns are discontinuous*, Journal of financial economics, 3 (1976), pp. 125–144.
- [25] C. MILLER AND I. YANG, *Optimal control of conditional value-at-risk in continuous time*, SIAM Journal on Control and Optimization, 55 (2017), pp. 856–884.
- [26] A. M. OBERMAN, *Convergent difference schemes for degenerate elliptic and parabolic equations: Hamilton–jacobi equations and free boundary problems*, SIAM Journal on Numerical Analysis, 44 (2006), pp. 879–895.
- [27] R. ROCKAFELLAR AND S. URYASEV, *Optimization of conditional value-at-risk*, Journal of Risk, 2 (2000), pp. 21–41.
- [28] R. T. ROCKAFELLAR, S. URYASEV, ET AL., *Optimization of conditional value-at-risk*, Journal of risk, 2 (2000), pp. 21–42.
- [29] M. J. RUIJTER AND C. W. OOSTERLEE, *Two-dimensional Fourier cosine series expansion method for pricing financial options*, SIAM Journal on Scientific Computing, 34 (2012), pp. B642–B671.
- [30] E. M. STEIN AND R. SHAKARCHI, *Fourier analysis: an introduction*, vol. 1, Princeton University Press, 2011.
- [31] P. T. TRAN ET AL., *On the convergence proof of amsgrad and a new version*, IEEE Access, 7 (2019), pp. 61706–61716.
- [32] K. YOSIDA, *Functional analysis*, vol. 123, Springer Science & Business Media, 2012.
- [33] H. ZHANG AND D. DANG, *A monotone numerical integration method for mean-variance portfolio optimization under jump-diffusion models*, Mathematics and Computers in Simulation, 219 (2024), pp. 112–140.
- [34] H. ZHOU AND D.-M. DANG, *Numerical analysis of American option pricing in a two-asset jump-diffusion model*, Applied Numerical Mathematics, 216 (2025), pp. 98–126.