Passivity-Preserving Model Reduction for Transient Cross-Impact

Joshua Kyan Aalampour
Paris, France

March 14, 2026

Abstract

Transient cross-impact models driven by matrix-valued convolution kernels are expressive but often non-Markovian and high-dimensional. We observe that the admissibility condition on the kernel is equivalent to passivity of the associated causal operator, and identify a structurally important subclass, the Stieltjes propagators, for which this equivalence enables constructive model reduction. Discretizing the positive relaxation measure rather than approximating the kernel directly, we obtain finite-dimensional surrogates that inherit economic admissibility by construction. We prove quantitative bounds on the distance between the optimal execution strategies and costs under the full and reduced models, controlled by the L^1 kernel approximation error.

1 Introduction

Two disciplines, working half a century apart and unaware of each other, arrived at the same inequality. In systems theory it is called passivity: no input signal can extract energy from the system [10]. In market microstructure it is called admissibility: no trading program can extract profit from transient impact [5, 6]. The mathematics is identical. This paper takes the coincidence seriously, and uses it to construct reduced-order impact models that are economically admissible by design.

Transient cross-impact models describe how past trades in d assets affect current prices through a matrix-valued convolution kernel G [4]. Rich kernels capture long memory and cross-asset structure; they also make optimal execution a non-Markovian problem requiring integral-equation methods. The single-exponential model of Obizhaeva and Wang [7] trades expressiveness for tractability, collapsing a rich temporal structure into a single decay rate. The problem is to do better: reduce a general kernel to a finite-dimensional surrogate that is faithful enough to be useful and constrained enough to be honest.

The obstacle is subtle and structural. A kernel that admits no profitable round trips may, after naive truncation, become one that does. The approximation inherits the shape but loses the constraint, and the constraint was the whole point. Passivity-preserving model reduction [3] is the control-theoretic answer to the analogous engineering problem. Our contribution is to bring this machinery into the cross-impact setting, identify the kernel class where it applies cleanly, and prove that it works with quantitative error bounds.

The class we work with is the *Stieltjes propagators*: kernels of the form $G(t) = \int_{(0,\infty)} \mathrm{e}^{-\rho t}\, M(\,\mathrm{d}\rho)$, where M is a positive semidefinite matrix-valued measure on the decay-rate axis. This class contains the exponential kernels, their finite sums, and the power-law decays observed empirically [4], while excluding oscillatory patterns outside the completely monotone framework [2]. Its value lies in a structural trio: every Stieltjes propagator is automatically admissible (Proposition 2), admits a passive distributed realization indexed by decay rates (Proposition 3), and can be reduced by discretizing the relaxation measure rather than approximating the kernel directly (Theorem 1). The reduced kernel inherits admissibility by construction, not by verification. Theorem 2 then proves that the optimal execution strategy and cost under the reduced model converge to those under the full model, with the error controlled by $\|G - G_r\|_{L^1(0,T)}$ and constants depending on the coercivity margin η and the problem data.

We do not characterize all admissible kernels. Completely monotone representations [2], passivity theory [10], and structure-preserving model reduction [3] each have established literatures; the contribution here is the fusion. The cross-impact landscape is by now substantial. Tomas, Mastromatteo, and Benzaquen [9] introduced a first-principles approach to cross-impact model selection, and Rosenbaum and Tomas [8] characterized cross-impact kernels under both martingale-admissibility and no-statistical-arbitrage conditions. On the execution side, Abi Jaber and Neuman [1] solved the general propagator liquidation problem for singular kernels via infinite-dimensional stochastic control. Those papers either characterize admissible kernels or solve execution within a given propagator model; our contribution is orthogonal: a method for *reducing* the propagator itself while preserving the no-manipulation property by construction. The result is a finance-native theorem package that connects no-manipulation conditions [5, 6] to passive systems theory through a concrete kernel class, and delivers reduced models with structural guarantees and quantitative error control.

2 Setup and the admissibility equivalence

Consider an agent trading d assets over a fixed horizon $[0, T]$. The trading rate $u \in L^2([0, T]; \mathbb{R}^d)$ determines the inventory

$$x_t = x_0 - \int_0^t u_s \, \mathrm{d}s, \tag{1}$$

where $x_0 \in \mathbb{R}^d$ is the initial position. Past trades affect current prices through a matrix-valued kernel $G\colon [0, \infty) \to \mathbb{S}^d$, assumed locally integrable.

Definition 1 (Impact operator and cost functional)**.** The *causal impact operator* induced by G is

$$(\mathcal{K}_G u)(t) \;:=\; \int_0^t G(t-s)\, u_s \, \mathrm{d}s, \qquad \text{for a.e. } t \in [0, T]. \tag{2}$$

The *impact cost* of a trading program u is

$$C_G(u) \;:=\; \langle u, \mathcal{K}_G u \rangle_{L^2} \;=\; \int_0^T u_t^\top \, (\mathcal{K}_G u)(t) \, \mathrm{d}t. \tag{3}$$

Throughout, $\|G\|_{L^1(0,T)} := \int_0^T \|G(t)\| \, \mathrm{d}t$, where $\|\cdot\|$ denotes the spectral norm on $\mathbb{R}^{d\times d}$.

The symmetry $G(t) \in \mathbb{S}^d$ yields the double-integral representation

$$C_G(u) \;=\; \tfrac{1}{2} \int_0^T \!\! \int_0^T u_t^\top \, G_e(t-s) \, u_s \, \mathrm{d}s \, \mathrm{d}t, \tag{4}$$

where $G_e(t) := G(|t|)$ is the even extension of G to $\mathbb{R}$. The kernel G is *admissible* if $C_G(u) \geq 0$ for every $T > 0$ and every $u \in L^2([0, T]; \mathbb{R}^d)$: no trading program, on any horizon, can extract profit from transient impact alone. Since $C_G(u) = \langle u, \mathcal{K}_G u \rangle_{L^2}$, admissibility is precisely the condition that $\mathcal{K}_G$ is a *passive* operator in the sense of systems theory [10].

The next result characterizes admissibility through the spectral properties of the kernel.

Proposition 1 (Admissibility equivalence)**.** *For continuous $G \in L^1_{\mathrm{loc}}([0,\infty);\mathbb{S}^d)$ with $G(0)$ finite, the following are equivalent.*

(i) Admissibility. *$C_G(u) \geq 0$ for every $T > 0$ and every $u \in L^2([0,T];\mathbb{R}^d)$.*

(ii) Positive-definite kernel. *The even extension G_e is a positive-definite matrix-valued kernel on $\mathbb{R}$.*

When $G(0)$ is not finite (e.g. for power-law kernels), condition (ii) is not directly meaningful in the pointwise sense; admissibility must then be established by other means, such as the dissipative identity of Proposition 3.

Proof. (ii)⇒(i). If G_e is positive-definite, the right-hand side of (4) is non-negative for every u, so $C_G(u) \geq 0$.

(i)⇒(ii). This is a Bochner-type characterization of positive-definite matrix-valued kernels. The precise statement in the transient impact setting, formulated for discrete time grids and all portfolios, appears in Alfonsi, Klöck, and Schied [2]; the continuous-time L^2 formulation follows by standard approximation. See also [5, 6] for earlier versions under stronger regularity. □

Remark 1 (Passivity and round trips)*.* In systems theory [10], admissibility says that the causal operator $\mathcal{K}_G$, viewed as a map from input (trading rate) to output (price impact), is *passive*: no trading program can extract energy from the system. In particular, no round-trip strategy ($\int_0^T u_t \,\mathrm{d}t = 0$) is profitable, recovering the no-price-manipulation condition of [5, 6] as a special case. Making the passivity identification explicit is the starting point for importing passivity-preserving reduction methods into the cross-impact setting.

3 Stieltjes propagators

Proposition 1 characterizes admissibility through positive-definiteness of the even extension G_e. We now identify a class of kernels for which this property holds by construction and which additionally supports a passive state-space realization and a structure-preserving reduction.

Definition 2 (Stieltjes propagator)**.** A kernel $G\colon (0,\infty) \to \mathbb{S}^d$ is a *Stieltjes propagator* if there exists a positive $\mathbb{S}^d_+$-valued Borel measure M on $(0,\infty)$ such that

$$G(t) \;=\; \int_{(0,\infty)} \mathrm{e}^{-\rho t}\, M(\,\mathrm{d}\rho), \qquad t > 0, \tag{5}$$

and $G \in L^1([0,T];\mathbb{S}^d)$ for every $T > 0$. We write $\mathcal{G}_{\mathrm{St}}$ for this class and call M the *relaxation measure* of G.

The measure M encodes how cross-impact is distributed across decay rates: each $\rho > 0$ represents a time scale at which past trades dissipate, and M allocates the matrix-valued intensity of impact to each scale. Throughout, integrals against M are interpreted via a scalar control measure σ and density $A(\rho) = \mathrm{d}M/\,\mathrm{d}\sigma \in \mathbb{S}^d_+$, so that $\int f(\rho)\, M(\,\mathrm{d}\rho) = \int f(\rho)\, A(\rho)\, \sigma(\,\mathrm{d}\rho)$. The integrability condition ensures that the impact operator $\mathcal{K}_G$ is bounded on $L^2([0,T];\mathbb{R}^d)$.

Proposition 2 (Admissibility of Stieltjes propagators)**.** *Every $G \in \mathcal{G}_{\mathrm{St}}$ is admissible. When G_e is finite-valued on $\mathbb{R}$ (for instance if $\operatorname{tr} M((0,\infty)) < \infty$), G_e is a positive-definite matrix-valued function and admissibility follows from Proposition 1. In the general case, admissibility is established independently by the realization identity of Proposition 3.*

Proof. The even extension takes the form $G_e(t) = \int_{(0,\infty)} \mathrm{e}^{-\rho|t|}\, M(\,\mathrm{d}\rho)$. For each $\rho > 0$, the scalar function $t \mapsto \mathrm{e}^{-\rho|t|}$ is positive-definite on $\mathbb{R}$. For any $(t_1, v_1), \dots, (t_n, v_n)$ in $\mathbb{R} \times \mathbb{R}^d$,

$$\sum_{i,j=1}^{n} v_i^\top G_e(t_i - t_j)\, v_j \;=\; \int_{(0,\infty)} \sum_{i,j=1}^{n} \mathrm{e}^{-\rho|t_i - t_j|}\, v_i^\top M(\,\mathrm{d}\rho)\, v_j. \tag{6}$$

The integrand is non-negative for each ρ. To see this, decompose M locally as $A(\rho)\,\sigma(\,\mathrm{d}\rho)$ with $A(\rho) \in \mathbb{S}_+^d$ and σ a positive scalar measure. Writing $A(\rho) = L(\rho)^\top L(\rho)$ and setting $w_i = L(\rho)\, v_i$, the sum reduces to $\sum_{i,j} \mathrm{e}^{-\rho|t_i - t_j|}\, w_i^\top w_j \geq 0$ by the scalar positive-definiteness of $\mathrm{e}^{-\rho|\cdot|}$. When $G_e(0)$ is finite, this shows that G_e is positive-definite in the classical sense, and admissibility follows from Proposition 1. For the general case, including singular kernels where $G_e(0)$ may be infinite, admissibility is proved directly by the dissipative identity (11) in Proposition 3 below, which yields $C_G(u) \geq 0$ without reference to pointwise kernel evaluation. □

Remark 2 (Scope of the class). In the scalar case, $\mathcal{G}_{\mathrm{St}}$ coincides with the locally integrable completely monotone functions on $(0, \infty)$ that vanish at infinity (equivalently, whose Bernstein representing measure has no atom at zero), characterized as Laplace transforms of positive measures on $(0, \infty)$; see [2] for the matrix-valued extension. Constant kernels (permanent impact) are excluded because Definition 2 places M on $(0, \infty)$, not $[0, \infty)$; this is by design, as the paper concerns transient impact. The class includes the single-exponential kernel of Obizhaeva and Wang [7], finite sums of exponentials, and power-law kernels $G(t) = c\,t^{-\alpha}$ with $\alpha \in (0, 1)$, which appear widely in empirical studies of market impact [4]. Not every admissible kernel is a Stieltjes propagator; oscillatory or non-monotone decay patterns fall outside $\mathcal{G}_{\mathrm{St}}$. For singular kernels such as power laws, where $G(0)$ is not finite, the pointwise equivalence of Proposition 1 does not apply; admissibility is instead established through the dissipative realization of Proposition 3, where the decomposition (11) shows $C_G(u) \geq 0$ without reference to G_e as a pointwise-defined kernel. The value of this class lies not in exhaustive generality but in the structural trio it supports: admissibility, passive realization, and structure-preserving reduction.

4 Passive realization

Every Stieltjes propagator admits a state-space realization indexed by the continuum of decay rates in its relaxation measure. The construction is reminiscent of a spectral decomposition in physics: a single complex object (the kernel) is resolved into a family of simple modes (exponential decays), each evolving independently, and the aggregate behavior emerges from their superposition. The realization carries a natural storage function whose dissipative structure makes passivity not merely provable but visible.

Let $G \in \mathcal{G}_{\mathrm{St}}$ with relaxation measure M. For a trading rate $u \in L^2([0, T]; \mathbb{R}^d)$, define the *distributed state* $z_t(\rho) \in \mathbb{R}^d$ by the family of ODEs

$$\partial_t z_t(\rho) = -\rho\, z_t(\rho) + u_t, \qquad z_0(\rho) = 0, \qquad \rho \in (0, \infty), \tag{7}$$

and the *output*

$$y_t \;=\; \int_{(0,\infty)} M(\,\mathrm{d}\rho)\, z_t(\rho). \tag{8}$$

Proposition 3 (Passive distributed realization). *Let $G \in \mathcal{G}_{\mathrm{St}}$ with relaxation measure M, and let $z_t(\rho)$ and y_t be defined by (7)–(8). Then $y_t = (\mathcal{K}_G u)(t)$ for almost every $t \in [0, T]$. Define the* *storage function*

$$S(t) \;:=\; \tfrac{1}{2}\int_{(0,\infty)} z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, z_t(\rho). \tag{9}$$

When M is a finite measure, S is absolutely continuous and satisfies, for almost every $t \in [0,T]$, the dissipative identity

$$\frac{\mathrm{d}S}{\mathrm{d}t} \;=\; u_t^\top\, y_t \;-\; \int_{(0,\infty)} \rho\, z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, z_t(\rho). \tag{10}$$

For general M, the integrated energy identity

$$\langle u,\, \mathcal{K}_G u\rangle_{L^2} \;=\; S(T) \;+\; \int_0^T\!\!\int_{(0,\infty)} \rho\, z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, z_t(\rho)\, \mathrm{d}t \;\geq\; 0 \tag{11}$$

holds, and in particular $\mathcal{K}_G$ is passive: $\langle u,\, \mathcal{K}_G u\rangle_{L^2} \geq 0$ for every $u \in L^2([0,T];\mathbb{R}^d)$.

Proof. The state equation (7) has the explicit solution $z_t(\rho) = \int_0^t \mathrm{e}^{-\rho(t-s)}\, u_s\, \mathrm{d}s$, so

$$y_t \;=\; \int_{(0,\infty)} M(\,\mathrm{d}\rho) \int_0^t \mathrm{e}^{-\rho(t-s)}\, u_s\, \mathrm{d}s \;=\; \int_0^t G(t-s)\, u_s\, \mathrm{d}s \;=\; (\mathcal{K}_G u)(t).$$

For the dissipative identity, differentiate S using (7) and the symmetry of M:

$$\begin{aligned}
\frac{\mathrm{d}S}{\mathrm{d}t} \;&=\; \int_{(0,\infty)} z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, \big[-\rho\, z_t(\rho) + u_t\big] \\
&=\; -\int_{(0,\infty)} \rho\, z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, z_t(\rho) \;+\; \left[\int_{(0,\infty)} M(\,\mathrm{d}\rho)\, z_t(\rho)\right]^\top u_t \\
&=\; -\int_{(0,\infty)} \rho\, z_t(\rho)^\top\, M(\,\mathrm{d}\rho)\, z_t(\rho) \;+\; u_t^\top\, y_t.
\end{aligned} \tag{12}$$

The exchange of differentiation and integration is justified when M is finite: the explicit solution $z_t(\rho) = \int_0^t \mathrm{e}^{-\rho(t-s)}\, u_s\, \mathrm{d}s$ satisfies $\|z_t(\rho)\| \leq \sqrt{T}\, \|u\|_{L^2}$ uniformly in ρ, controlling the cross term against any finite M, and $\rho\, \|z_t(\rho)\|^2 \leq \frac{1}{2}\|u\|_{L^2}^2$, controlling the ρ-weighted dissipation term. When M is finite, integrating (12) over $[0,T]$ and using $S(0) = 0$ yields (11) directly. For general M, let $M_N := M|_{[1/N,N]}$ and let G_N be the induced kernel. Since $0 \preceq G_N(t) \preceq G(t)$ and $G \in L^1(0,T)$, dominated convergence gives $\|G_N - G\|_{L^1(0,T)} \to 0$, and Young's inequality then yields $\langle u,\, \mathcal{K}_{G_N} u\rangle_{L^2} \to \langle u,\, \mathcal{K}_G u\rangle_{L^2}$. Writing $S_N(T)$ and D_N for the storage and dissipation terms under M_N, both are non-negative and increase with N, so monotone convergence gives (11) in the limit. Passivity follows. □

Remark 3 (Interpretation)*.* Identity (11) gives the impact cost a physical grammar. The storage $S(T)$ is potential energy: impact that the market has absorbed but not yet forgotten, latent reversion waiting to act on future prices. The dissipation integral is thermal loss: energy that has already decayed into the microstructure and cannot be recovered. Fast modes burn hot and die quickly; slow modes smolder. A round-trip strategy that appears costless in a naive model is, through this lens, one whose energy has been reclassified (stored or dissipated) but never destroyed. This realization is infinite-dimensional, indexed by the continuum of decay rates in the support of M. The next section replaces it with a finite-dimensional surrogate.

5 Arbitrage-preserving reduction

The distributed realization of Section 4 is indexed by the full support of the relaxation measure M, which may be a continuum. We now replace M by a finitely supported positive measure, producing a reduced kernel that inherits admissibility by construction.

Fix $r \geq 1$ decay rates $\rho_1, \dots, \rho_r > 0$ and residue matrices $A_1, \dots, A_r \in \mathbb{S}^d_+$. Define the atomic measure

$$M_r \;:=\; \sum_{k=1}^{r} A_k\, \delta_{\rho_k} \tag{13}$$

and the reduced kernel

$$G_r(t) \;:=\; \int_{(0,\infty)} \mathrm{e}^{-\rho t}\, M_r(\,\mathrm{d}\rho) \;=\; \sum_{k=1}^{r} A_k\, \mathrm{e}^{-\rho_k t}, \qquad t > 0. \tag{14}$$

Theorem 1 (Arbitrage-preserving reduction)**.** *Let $G \in \mathcal{G}_{\mathrm{St}}$, and let G_r be defined by (13)–(14) for any choice of nodes $\rho_k > 0$ and residues $A_k \in \mathbb{S}^d_+$. Then:*

(i) *$G_r \in \mathcal{G}_{\mathrm{St}}$, and in particular G_r is admissible.*

(ii) *The impact operator $\mathcal{K}_{G_r}$ admits the finite-dimensional passive realization*

$$\dot{z}^{(k)}_t = -\rho_k\, z^{(k)}_t + u_t, \qquad z^{(k)}_0 = 0, \qquad k = 1, \dots, r, \tag{15}$$

$$(\mathcal{K}_{G_r} u)(t) \;=\; \sum_{k=1}^{r} A_k\, z^{(k)}_t, \tag{16}$$

with storage function

$$S_r(t) \;=\; \tfrac{1}{2} \sum_{k=1}^{r} \big(z^{(k)}_t\big)^{\top} A_k\, z^{(k)}_t. \tag{17}$$

(iii) Approximation density. *For every $\varepsilon > 0$, there exist $r \geq 1$, nodes $\rho_k > 0$, and residues $A_k \in \mathbb{S}^d_+$ such that $\|G - G_r\|_{L^1(0,T)} < \varepsilon$.*

Proof. Parts (i) and (ii). The measure M_r is a finite sum of Dirac masses weighted by matrices in $\mathbb{S}^d_+$, hence a positive $\mathbb{S}^d_+$-valued Borel measure on $(0, \infty)$. By Definition 2, $G_r \in \mathcal{G}_{\mathrm{St}}$, and admissibility follows from Proposition 2. The realization (15)–(16) and storage function (17) are the specialization of Proposition 3 to the atomic measure M_r.

Part (iii). Fix $\varepsilon > 0$. Restrict M to $[1/N, N]$ for N large enough that $\|G - G_N\|_{L^1(0,T)} < \varepsilon/2$, where G_N is the kernel induced by the restricted measure; this holds by dominated convergence since $G \in L^1(0,T)$. Partition $[1/N, N]$ into subintervals $I_1, \dots, I_r$, choose a node $\rho_k \in I_k$ for each k, and set $A_k := \int_{I_k} M(\,\mathrm{d}\rho) \in \mathbb{S}^d_+$. Then

$$G_N(t) - G_r(t) \;=\; \sum_{k=1}^{r} \int_{I_k} \big(\mathrm{e}^{-\rho t} - \mathrm{e}^{-\rho_k t}\big)\, M(\,\mathrm{d}\rho),$$

and uniform continuity of $\rho \mapsto \mathrm{e}^{-\rho t}$ on $[1/N, N] \times [0, T]$ ensures that the mesh can be chosen fine enough that $\|G_N - G_r\|_{L^1(0,T)} < \varepsilon/2$. The triangle inequality gives $\|G - G_r\|_{L^1(0,T)} < \varepsilon$. ☐

Remark 4 (Reducing the measure, not the kernel). The design principle is to approximate the relaxation measure M rather than the kernel G directly. This is the difference between sculpting within a block of marble and gluing fragments together afterward: any positive atomic measure M_r produces a reduced kernel $G_r \in \mathcal{G}_{\mathrm{St}}$ that is admissible by birth, not by repair. Generic kernel approximation [3] works on G and can destroy the positive-definiteness of G_e, reintroducing the very manipulation opportunities the model was built to exclude. The Stieltjes framework avoids this by never leaving the cone of positive measures. The proof of part (iii) is constructive: partition the support of M, assign each node to a subinterval, and integrate M over each bin to obtain the residues. The choice of nodes ρ_k and residues A_k determines the approximation quality $\|G - G_r\|_{L^1(0,T)}$; Section 6 quantifies the effect of this error on optimal execution.

6 Execution error bounds

We now quantify the effect of the kernel approximation on optimal execution. Consider the linear-quadratic objective

$$J_G(u) \;:=\; \tfrac{1}{2}\langle u,\, \mathcal{K}_G u\rangle_{L^2} \;+\; \tfrac{\eta}{2}\|u\|_{L^2}^2 \;+\; \tfrac{\lambda}{2}\int_0^T x_t^\top \,\Sigma\, x_t \,\mathrm{d}t, \tag{18}$$

where $\eta > 0$ is a temporary impact parameter, $\lambda \ge 0$ is an inventory risk penalty, $\Sigma \in \mathbb{S}_+^d$ is a covariance matrix, and $x_t = x_0 - \int_0^t u_s \,\mathrm{d}s$ is the inventory process. The three terms penalize, respectively, transient impact cost, trading speed, and exposure to inventory risk.

Since $\mathcal{K}_G$ is passive (Proposition 2) and $\eta > 0$, the functional J_G is η-strongly convex on $L^2([0,T];\mathbb{R}^d)$ and admits a unique minimizer $u^\star$. Write

$$\mathcal{A}_G\, u^\star = b, \tag{19}$$

where $\mathcal{A}_G$ is the self-adjoint, positive operator associated with the quadratic part of J_G and $b \in L^2([0,T];\mathbb{R}^d)$ encodes the linear dependence on the initial inventory x_0. Coercivity gives $\|\mathcal{A}_G^{-1}\|_{L^2\to L^2} \le \eta^{-1}$. Define the *optimal value* $V(G) := J_G(u^\star)$.

Replacing G by the reduced kernel G_r of Theorem 1 yields the reduced objective J_{G_r}, minimizer $u_r^\star$, operator $\mathcal{A}_{G_r}$, and value $V(G_r)$. The causal operator $\mathcal{K}_G$ is not self-adjoint. The self-adjoint operator associated with the quadratic form C_G is

$$(\mathcal{S}_G\, u)(t) \;:=\; \tfrac{1}{2}\int_0^T G_e(t-s)\, u_s \,\mathrm{d}s, \tag{20}$$

which satisfies $\langle u,\, \mathcal{S}_G\, u\rangle_{L^2} = C_G(u)$ by (4). Since η and the inventory penalty do not depend on the kernel, $\mathcal{A}_G - \mathcal{A}_{G_r} = \mathcal{S}_G - \mathcal{S}_{G_r}$.

Lemma 1 (Operator perturbation bound)**.** *For any kernels G, $G_r \in L^1([0,T];\mathbb{S}^d)$,*

$$\|\mathcal{S}_G - \mathcal{S}_{G_r}\|_{L^2\to L^2} \;\le\; \|G - G_r\|_{L^1(0,T)}. \tag{21}$$

Proof. The operator $\mathcal{S}_G - \mathcal{S}_{G_r}$ acts by convolution with $\tfrac{1}{2}(G_e - G_{r,e})$. Young's inequality gives the bound $\tfrac{1}{2}\|G_e - G_{r,e}\|_{L^1(-T,T)}$, which equals $\|G - G_r\|_{L^1(0,T)}$ by evenness. □

Theorem 2 (Value and optimizer stability)**.** *Let $G \in \mathcal{G}_{\mathrm{St}}$ and let G_r be a reduced kernel as in Theorem 1. Assume $\eta > 0$. Then*

$$\|u_r^\star - u^\star\|_{L^2} \;\le\; \frac{\|u^\star\|_{L^2}}{\eta}\,\|G - G_r\|_{L^1(0,T)}, \tag{22}$$

$$|V(G_r) - V(G)| \;\le\; \frac{\|b\|_{L^2}\,\|u^\star\|_{L^2}}{2\eta}\,\|G - G_r\|_{L^1(0,T)}. \tag{23}$$

Proof. Operator bound. Since $\mathcal{A}_G - \mathcal{A}_{G_r} = \mathcal{S}_G - \mathcal{S}_{G_r}$, Lemma 1 gives

$$\|\mathcal{A}_G - \mathcal{A}_{G_r}\|_{L^2\to L^2} \;\le\; \|G - G_r\|_{L^1(0,T)}. \tag{24}$$

Optimizer bound. Both minimizers satisfy the same structural equation with the same right-hand side: $\mathcal{A}_G\, u^\star = b = \mathcal{A}_{G_r}\, u_r^\star$. The resolvent identity gives

$$u_r^\star - u^\star \;=\; \mathcal{A}_{G_r}^{-1}(\mathcal{A}_G - \mathcal{A}_{G_r})\, u^\star,$$

so

$$\|u_r^\star - u^\star\|_{L^2} \;\le\; \|\mathcal{A}_{G_r}^{-1}\|\,\|\mathcal{A}_G - \mathcal{A}_{G_r}\|\,\|u^\star\|_{L^2} \;\le\; \frac{\|u^\star\|_{L^2}}{\eta}\,\|G - G_r\|_{L^1(0,T)}.$$

Value bound. For any quadratic functional $J(u) = \frac{1}{2}\langle u, Au\rangle_{L^2} - \langle b, u\rangle_{L^2} + c_0$ with minimizer $Au^\star = b$, the optimal value is $V = -\frac{1}{2}\langle b, u^\star\rangle_{L^2} + c_0$. Since the constant c_0 does not depend on the kernel,

$$V(G_r) - V(G) \;=\; \tfrac{1}{2}\langle b,\, u^\star - u_r^\star\rangle_{L^2}.$$

The Cauchy–Schwarz inequality and (22) yield (23). □

Remark 5 (Role of the constants)*.* The constants in Theorem 2 depend on the coercivity margin η and the problem scale through $\|u^\star\|_{L^2}$ and $\|b\|_{L^2}$. Using $\|u^\star\|_{L^2} \le \eta^{-1}\|b\|_{L^2}$, the bounds simplify to $O\big(\eta^{-2}\|b\|_{L^2}\,\|G - G_r\|_{L^1}\big)$ and $O\big(\eta^{-2}\|b\|_{L^2}^2\,\|G - G_r\|_{L^1}\big)$ respectively, isolating the sensitivity to the temporary impact parameter. The assumption $\eta > 0$ is not a technicality: without temporary impact, the objective may fail to be coercive, and the perturbation theory breaks down. A hard liquidation constraint $x_T = 0$ or a terminal penalty can be incorporated without changing the analysis, since neither depends on the kernel.

7 Conclusion

A market that cannot be manipulated and a circuit that cannot generate energy satisfy the same inequality. The cost functional identifies them: $C_G(u) = \langle u,\, \mathcal{K}_G u\rangle_{L^2} \ge 0$ is simultaneously the definition of admissibility and the definition of passivity. Proposition 1 then characterizes the condition through the kernel. Once this is recognized, fifty years of passivity-preserving engineering become available to the financial modeler, and the question shifts from *whether* reduced models can preserve admissibility to *how cheaply*.

We chose to work with Stieltjes propagators because they are where the structure concentrates. The class is rich enough to contain the empirically dominant kernel shapes [4, 7] and rigid enough that reduction cannot destroy what matters. Within the class of finance-native propagator models built from completely monotone decay structures, we are not aware of a broader subclass that retains all three properties simultaneously: automatic admissibility, passive realization, and structure-preserving reduction. In the control-theoretic literature, the positive-real class is broader and admits

its own passivity-preserving reduction methods [3]; our contribution is the specialization to transient impact, where the relaxation-measure discretization provides a particularly transparent reduction mechanism.

Several directions remain open, and they pull toward different disciplines. Nonlinear dissipative systems in the sense of Willems [10] generalize passivity beyond linear operators; whether the corresponding nonlinear impact models retain tractable reduction is a question that belongs as much to nonlinear control as to finance. Multi-agent settings introduce strategic interaction, but the individual passivity constraints survive and may discipline the equilibrium, much as conservation laws constrain many-body systems even when individual trajectories become intractable. On the empirical side, calibrating the relaxation measure M from propagator data is an inverse Laplace problem constrained to the positive cone: a well-studied problem with known regularization techniques, waiting for market data to meet it. Finally, the overlap between Stieltjes propagators and affine Volterra processes suggests that passivity-preserving reduction may extend to stochastic impact models. The bridge between the deterministic kernel theory developed here and the stochastic world is narrow but visible.

References

[1] Eduardo Abi Jaber and Eyal Neuman. Optimal liquidation with signals: the general propagator case. *Mathematical Finance*, 35(4):841–866, 2025.

[2] Aurélien Alfonsi, Florian Klöck, and Alexander Schied. Multivariate transient price impact and matrix-valued positive definite functions. *Mathematics of Operations Research*, 41(3):914–934, 2016.

[3] Athanasios C. Antoulas. *Approximation of Large-Scale Dynamical Systems*. SIAM, Philadelphia, 2005.

[4] Jean-Philippe Bouchaud, J. Doyne Farmer, and Fabrizio Lillo. How markets slowly digest changes in supply and demand. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editors, *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–160. North-Holland, 2009.

[5] Jim Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759, 2010.

[6] Gur Huberman and Werner Stanzl. Price manipulation and quasi-arbitrage. *Econometrica*, 72 (4):1247–1275, 2004.

[7] Anna A. Obizhaeva and Jiang Wang. Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1):1–32, 2013.

[8] Mathieu Rosenbaum and Mehdi Tomas. A characterisation of cross-impact kernels. *Frontiers of Mathematical Finance*, 1(4):491–523, 2022.

[9] Mehdi Tomas, Iacopo Mastromatteo, and Michael Benzaquen. How to build a cross-impact model from first principles: theoretical requirements and empirical results. *Quantitative Finance*, 22(6):1017–1036, 2022.

[10] Jan C. Willems. Dissipative dynamical systems, part I: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.