

Relativistic Asset Pricing: A Formal Treatment

Author: Daniel Cheah **Affiliation:** Space Finance Institute (independent research initiative) **Contact:**

danielcheah@spacefi.institute **Date:** July 2026 — Version 2.0, consolidated (revises Version 1, 1 November 2025, SSRN 5693543, and consolidates the companion technical note, SSRN 5721142)

License: CC BY-NC-ND 4.0 This work is distributed under the Creative Commons Attribution–NonCommercial–NoDerivatives 4.0 International License (CC BY-NC-ND 4.0), allowing free sharing with attribution but prohibiting commercial use or modification.

WORKING PAPER — COMMENTS WELCOME

Abstract

This paper develops **Relativistic Asset Pricing (RAP)**, a framework for financial markets constrained by the finite speed of information.

Where classical models assume instantaneous diffusion, RAP imposes an information-velocity limit $c_{\text{eff}} \leq c$, defining causal domains analogous to light cones in spacetime. Arbitrage opportunities become realisable only within overlapping informational cones of market participants.

We restate the no-arbitrage condition under causal filtrations $F_{\Lambda(t,x)}$: the sufficiency direction follows from the classical fundamental theorem by a conditioning argument, while the converse belongs to the established theory of asset pricing under delayed and restricted information (Kabanov and Stricker 2006, and continuous-time extensions). We also prove an obstruction result: hyperbolic (wave-type) dynamics cannot model the traded price itself, because solutions of second-order-in-time equations have finite-variation paths and therefore carry no martingale part for any change of measure to act on — finite-speed causality must enter through delayed filtrations or through the covariance of the driving noise, not through wave dynamics of the price.

RAP's constructive contribution is the geometry: delays are not exogenous constants but $\tau_{ij} = d_{ij}/c_{\text{eff}}$ on a spatial venue graph, which makes the bound on arbitrage reach quantitative. We derive the **maximum causal arbitrage distance**

$$D_{\text{max}} = \min \left\{ c \Delta t^*, c_{\text{eff}} \Delta t^* + c_{\text{eff}} \sigma_{\tau} \Phi^{-1}(1 - \theta) \right\}$$


from the survival probability of an exploitable signal over stochastic transit, where Δt^* is signal persistence, σ_{τ} latency jitter, and θ the required success probability. Demanding reliability ($\theta > 1/2$) *shrinks* reach below the mean-transit frontier; speculative strategies ($\theta < 1/2$) extend it, capped by the vacuum light cone. We further derive the delayed-hedging cost of communication blackouts, which is maturity-dependent: decisive for short-dated, high-gamma claims and basis points for long-dated ones, with a worked Earth–Mars example.

The model reconciles network physics and financial efficiency: as $c_{\text{eff}} \rightarrow \infty$ it converges to classical finance, while finite c_{eff} and jitter capture real-world latency-bound arbitrage. Empirical analogues are drawn from high-frequency trading networks and interplanetary communication delays. RAP reframes market efficiency as a thermodynamic problem: synchrony requires energy, bandwidth and infrastructure. The same dynamics inform market-design policy on the physical constraints of fairness and latency regulation.

This version consolidates the former companion technical note (SSRN 5721142) into a single paper; that record now carries a supersession notice and erratum.

RELATIVISTIC ASSET PRICING: A Formal Treatment

Information propagation at finite speed bounds arbitrage opportunities



Daniel Cheah — Space Finance Institute | Companion technical note: SSRN 5721142

Graphical Abstract — Information propagation at finite speed constrains arbitrage and defines causal market domains.

SSRN Metadata

Keywords: relativistic finance; causal filtration; latency arbitrage; delayed information; econophysics; information geometry; market microstructure; high-frequency trading **JEL Codes:** G12, G14, C02, C60 **Research Fields:** Quantitative Finance; Financial Economics; Econophysics; Computational Finance

Notation (Summary)

Symbol	Definition	Units
c	Vacuum speed of light (hard causal bound)	[L/T]
c_{eff}	Effective propagation velocity of market data (fibre $\approx 0.67c$, microwave $\approx c$)	[L/T]
$\mathbf{M}$	Market manifold ($\mathbb{R}^3 \times \mathbb{R}$) with coordinates (ct, x^1, x^2, x^3)	—
$\mathbf{g}_{\{\mu\nu\}}$	Informational metric (Minkowski-like)	—
$\Lambda(\mathbf{t}, \mathbf{x})$	Past light cone (causal domain) of $(\mathbf{t}, \mathbf{x})$	—
$\mathbf{F}_{\{\Lambda(\mathbf{t}, \mathbf{x})\}}$	Causal filtration (σ -algebra restricted to Λ)	—
σ_{τ}	Standard deviation of network latency jitter	[T]
σ	Causal smoothing length, $\sigma = c_{\text{eff}} \cdot \sigma_{\tau}$	[L]
θ	Required probability that a signal survives transit (net of costs)	[dimensionless]
$\tau(\mathbf{x}, \mathbf{x}')$	Communication delay =	$\mathbf{x} - \mathbf{x}'$
Δt^*	Decay time of exploitable price differentials	[T]
$D_{\max}$	Maximum causal arbitrage distance	[L]
Γ	Option gamma	[1/price]
τ_{rt}	Round-trip communication blackout, $2d/c$	[T]

1 Introduction and Motivation

Classical finance assumes instantaneous information transmission, implying infinite-speed markets.

Relativistic Asset Pricing (RAP) imposes a finite information velocity, bounding arbitrage opportunities within causal domains defined by light-cone-like boundaries. Arbitrage becomes realisable only within overlapping informational cones of traders.

We note that special relativity encompasses multiple phenomena: time dilation due to relative motion (Lorentz factor γ), and finite signal propagation. RAP addresses the latter, information-velocity constraints, which dominate at terrestrial velocities where $v/c \approx 10^{-8}$ and time dilation is negligible ($\gamma \approx 1 + 10^{-16}$).

Relation to prior work. The technical lineage is acknowledged explicitly, and its oldest entries precede the microstructure era: Krugman (1978, published 2010) derived interest parity for goods in relativistic transit, and Haug (2004) introduced relativistic corrections to volatility and rates for practitioners. The fundamental theorem of asset pricing under *delayed and restricted information* is established: Kabanov and Stricker (2006) prove the Dalang–Morton–Willinger theorem in that setting, with continuous-time, two-filtration extensions in the subsequent literature. Finite-velocity price processes have a developed option-pricing theory via telegraph processes (Kolesnik and Ratanov 2013; Ratanov 2007) and the relativistic heat-equation route (Trzetrzelewski 2017), whose prices converge to Black–Scholes as $c \rightarrow \infty$. The spatial geometry of latency arbitrage was quantified by Wissner-Gross and Freer (2010), who located optimal intermediate nodes for trading spacelike-separated venues; the regulatory consequences of light-speed limits were mapped by Angel (2014); latency as an economically informative variable is documented by Kirilenko and Lamacie (2015); order transmission along communication lines with relativistic delay is modelled with queueing theory by Lerner (2021); relativistic option pricing across reference frames is developed by Carvalho and Gaspar (2021); and the physical impossibility of perfect informational efficiency under light-cone constraints is argued by Petrov (2025). RAP's contribution relative to this literature is integrative and geometric: per-location causal filtrations indexed by a spatial venue graph, an obstruction result delimiting which dynamics can carry martingale prices (Section 4), an operational and calibratable bound on arbitrage reach (Section 6), a derived delayed-hedging cost for communication blackouts (Section 7), and the market-design consequences (Section 8).

2 Information Geometry of Markets

Market spacetime: $M \cong \mathbb{R}^3 \times \mathbb{R}$ with coordinates (ct, x^1, x^2, x^3) , where $\mathbf{c_eff}$ represents the information speed limit (the effective propagation velocity of market data through communication networks) and $\mathbf{c}$ the vacuum speed of light, the hard causal bound no signal evades.[^satellite-time-dilation]

The informational metric adopts Minkowski form $ds^2 = c^2 dt^2 - |dx|^2$.

Assumption (Informational Causality Metric): This is an informational geometry; $ds^2 \geq 0$ represents feasible communication, not physical spacetime.

Each observer defines a causal filtration

$$F_{\Lambda(t,x)} = \sigma\{S_{t',x'} : (t', x') \in \Lambda(t, x)\}$$

limiting access to information. For the past light cone we write $\Lambda(t, x) = \{(s, y) : s \leq t, \|x - y\| \leq c(t - s)\}$, so that $F_{\Lambda(t,x)} \subseteq F_t$, the full filtration.

While Bloss and Kleinknecht (2016) and Carvalho and Gaspar (2021) explore time dilation effects on option pricing for market participants in different relativistic reference frames (particularly relevant for interplanetary trading), RAP focuses on information propagation delays within causal light cones, a phenomenon directly observable in terrestrial high-frequency trading networks.

[^satellite-time-dilation]: For satellite-based market infrastructure, both information delay ($\tau = d/c_{\text{eff}}$) and time dilation (proper time corrections) could theoretically compound. GPS already corrects for ~ 38 microseconds/day of relativistic time dilation. However, such corrections are several orders of magnitude below the millisecond-scale latencies that drive HFT arbitrage.

3 Causality-Bounded Arbitrage

Theorem (No-arbitrage under causal filtrations — restated). Let markets on $(M, g_{\{\mu\nu\}})$ carry locally traded assets $S(x)$ with money account B , and let admissible strategies at (t, x) be predictable and adapted to $F_{\Lambda(t, x)}$.

(a) Sufficiency. Suppose there exists an equivalent measure Q under which each discounted price $S(x)/B$ is a martingale for the full filtration $\{F_t\}$, and local prices are locally observable ($S_t(x)/B_t$ is $F_{\Lambda(t, x)}$ -measurable). Then for all $T > t$,

$$E_Q[S_T(x)/B_T \mid F_{\Lambda(t, x)}] = S_t(x)/B_t,$$

by the tower property, since $F_{\Lambda(t, x)} \subseteq F_t$; and no admissible causal strategy admits arbitrage, since every F_{Λ} -adapted strategy is also F_t -adapted and the classical theorem applies (Harrison and Kreps 1979; Delbaen and Schachermayer 1994).

(b) Converse. The existence of an equivalent martingale measure given only the *absence of causal arbitrage* is the fundamental theorem of asset pricing under delayed and restricted information. In discrete time this is Kabanov and Stricker (2006); continuous-time two-filtration versions exist in the subsequent literature. RAP specialises that theory: the information lag of an agent at y with respect to venue x is not an exogenous constant but the light-time $\tau(x, y) = \|x - y\|/c_{\text{eff}}$, so the delayed-information structure is generated by the spatial geometry of the market itself.

As $c_{\text{eff}} \rightarrow \infty$, Λ expands to M , the causal filtration converges to the full filtration, and the statement reduces to the classical FTAP.

Remark (what is and is not claimed). Direction (a) is a corollary of the classical theorem; direction (b) is cited, not re-proved. The framework's original content lies in the geometric specialisation, the obstruction result of Section 4, and the operational bound of Section 6. Version 1 of this paper presented the statement as a new theorem ("CBPT") with a separation-argument sketch; that presentation overstated novelty and understated the delicacy of the converse, and is superseded by the present statement.

4 The Obstruction: Why Hyperbolic Price Dynamics Cannot Support Martingale Pricing

A tempting route to finite-speed markets — and the one taken in the superseded technical note (SSRN 5721142, v1) — is to model the log-price itself with a damped-wave (hyperbolic) stochastic PDE,

$$\partial_{tt} \delta + 2\gamma \partial_t \delta = c^2 \Delta \delta + \sigma_S \dot{W}, \tag{1}$$

and to seek a change of measure making $e^{-rt}P$, $P = e^\delta$, a martingale under the causal filtration. That route is closed, for a structural reason worth recording:

Proposition 1 (Obstruction). Recast (1) as the first-order system $d\delta = v dt$, $dv = (c^2 \Delta \delta - 2\gamma v) dt + \sigma_S dW$. Suppose the noise is spatially regular enough that $v(t, x)$ is a genuine function-valued process (as under σ -mollified noise, the specification needed for calibration and numerical stability). Then for each x , $\delta(t, x) = \delta(0, x) + \int_0^t v(s, x) ds$ has absolutely continuous paths in t , hence zero quadratic variation. Consequently $P = e^\delta$ satisfies $dP = P v dt$: the price is a finite-variation process with **no martingale part**, and $e^{-rt}P$ is a local martingale only if $v \equiv r$.

identically, which means deterministic prices. No equivalent change of measure alters this: Girsanov shifts the drift of processes possessing a Brownian part; it cannot create a martingale part where none exists. $\square$

Remark 1. The noise in (1) enters the *velocity* $v = \partial_t \delta$, not the price level; that is what makes the dynamics hyperbolic, and it is also what removes the price's quadratic variation. Martingale prices require noise at the level of the price increment.

Remark 2 (the unsmoothed case). With genuine space-time white noise in $d = 3$, the situation is worse rather than better: the stochastic wave equation has no function-valued solution (the solution is a distribution in x), so $P = e^\delta$ is not even defined pointwise. Either way, no theorem of the intended form exists.

Remark 3 (where causality must live instead). Two consistent routes remain, and this paper uses the first:

1. **Delayed filtrations.** Prices follow ordinary semimartingale dynamics at each venue; finite-speed causality enters through *who can know what, when* (the causal filtrations of Section 2) and through delayed observation of remote venues. This is the setting of the FTAP under delayed and restricted information invoked in Section 3.
2. **Causal covariance.** Prices are driven by martingale noise fields that are white in time at each location but spatially correlated only within light cones (cone-shaped "ambit sets" in the sense of the ambit-stochastics literature of Barndorff-Nielsen, Benth and Veraart). Development of this route, including its testable implication that cross-venue return correlations ramp up over sampling horizons of the order of inter-venue light time (an Epps-type effect), is left to future work.

Hyperbolic equations retain a legitimate role as models of *information transport* — the micro-foundation for finite c_{eff} is a persistent random walk whose density satisfies a telegraph equation (Appendix A) — and as simulation devices for information fields (field simulator in the SSRN ancillary notebook). They are withdrawn as models of the traded price.

5 Delayed Observation and Causal Correlation

Causal correlation kernel (corrected form). Cross-location correlation of price-relevant signals is modelled as Gaussian in the dispersion of *effective* transit around its mean, with hard support inside the vacuum cone:

$$f(|x - x'|, |t - t'|; c_{\text{eff}}, c) = \begin{cases} \exp\left(-\frac{(|x - x'| - c_{\text{eff}} |t - t'|)^2}{\sigma^2}\right), & |x - x'| \leq c |t - t'|, \\ 0, & \text{otherwise.} \end{cases}$$

Here $\sigma = c_{\text{eff}} \sigma_\tau$ is the *causal smoothing length* derived from stochastic latency jitter: signal transmission time is $\tau = |x - x'|/c_{\text{eff}} + \varepsilon$ with $\varepsilon \sim N(0, \sigma_\tau^2)$, reflecting routing, congestion and processing variability around the mean propagation speed. Lerner (2021) supplies a queueing-theoretic microfoundation for such distortion: nodes along the transmission path trade against the passing signal, deforming it even at zero reaction time. Because $c_{\text{eff}} < c$, arrivals earlier than the mean are physically permitted, but never earlier than the vacuum light time $|x - x'|/c$. This is why the kernel's support is capped by the *vacuum* cone rather than the effective cone.

Typical HFT magnitudes: $\sigma \approx 2\text{--}30$ km (10–100 μs).

Assumption (Small Smoothing Regime): $\sigma \ll c_{\text{eff}} \Delta t^*$, ensuring the smoothing band is thin relative to market timescales.

The sharp and smoothed regimes compare as follows:

Regime	Sharp ($\sigma \rightarrow 0$)	Smooth (σ finite)
Physical basis	Perfect signals	Noisy networks

Regime	Sharp ($\sigma \rightarrow 0$)	Smooth (σ finite)
Kernel	$\delta(d - c_{\text{eff}} t)$	$\exp(-(d - c_{\text{eff}} t)^2 / \sigma^2)$, support $d \leq ct$
Causality	Exact	Exact (vacuum-cone capped)
Calibration	Impossible	Tractable
Simulations	Unstable	Stable
Applications	Theoretical	Empirical

Note. Stochastic reach beyond the *mean* frontier $c_{\text{eff}} \cdot \Delta t^*$ is legitimate (fast routing days), but nothing crosses the vacuum cone. Section 6 derives the reach bound from first principles rather than from this kernel; the correction history relative to Version 1 is recorded in the Version note.




Figure 1. Smoothed versus sharp causal domains. The solid line illustrates the mean propagation frontier ($c_{\text{eff}} t = |x - x'|$); the shaded halo represents stochastic dispersion of transit times with width σ . The halo is one-sided in physics: it may reach toward, but never beyond, the vacuum light cone (outer dashed boundary). As $\sigma \rightarrow 0$ the halo collapses to the mean frontier.

Delayed-observation chain rule (remark, not a new lemma). If the price observed at x' is the light-delayed remote price, $S_t^{\text{obs}}(x') := S_{t-\tau(x,x')}(x)$, then for $f(S,t)$ twice differentiable, classical Itô calculus applied to the time-shifted semimartingale gives

$$df(S_t^{\text{obs}}(x), t) = (\partial f / \partial t) dt + (\partial f / \partial S) dS_{t-\tau}(x) + \frac{1}{2} (\partial^2 f / \partial S^2) d\langle S \rangle_{t-\tau}(x).$$

This is a change of variables, not a new stochastic calculus; Version 1's label "relativistic Itô lemma" is retired. The rigorous home for dynamics with delays is the theory of stochastic functional differential equations (Mohammed 2013) and delayed Black–Scholes analysis (Arriojas, Hu, Mohammed and Pap 2007). What is specific to RAP is only that the delay is the light-time of the venue geometry.

6 The Maximum Distance Constraint (D_{max})

A profitable response at distance d to a signal observed at its source requires the signal to survive transit. Let Δt^* denote the decay time of exploitable price differentials, and let transit time be $d/c_{\text{eff}} + \varepsilon$ with $\varepsilon \sim N(0, \sigma_\tau^2)$, floored by the vacuum light time d/c . Let $\theta \in (0,1)$ be the required probability (net of costs) that the signal is still alive on arrival.

The strategy clears the reliability hurdle iff

$$\Pr(d/c_{\text{eff}} + \varepsilon \leq \Delta t^*) \geq \theta \iff d \leq c_{\text{eff}} \Delta t^* + c_{\text{eff}} \sigma_\tau \Phi^{-1}(1 - \theta),$$

subject to the vacuum bound $d \leq c \Delta t^*$. Hence

$$D_{\text{max}} = \min\{c \Delta t^*, c_{\text{eff}} \Delta t^* + c_{\text{eff}} \sigma_\tau \Phi^{-1}(1 - \theta)\}^{**}$$

Interpretation. The mean frontier is $c_{\text{eff}} \Delta t^*$. Jitter then cuts both ways, and the sign is governed by the reliability demanded:

- **Reliability regime ($\theta > 1/2$):** $\Phi^{-1}(1 - \theta) < 0$. Insisting on a high success probability *shrinks* reach below the mean frontier. Jitter is a tax on dependable strategies; reducing σ_τ recovers reach.
- **Speculative regime ($\theta < 1/2$):** $\Phi^{-1}(1 - \theta) > 0$. Strategies content with occasional success reach beyond the mean frontier on fast-routing draws, up to, and never beyond, the vacuum cone.

Dimensional check: $c_{\text{eff}} \Delta t^*$ and $c_{\text{eff}} \sigma_\tau$ share units [L]; $\Phi^{-1}(1 - \theta)$ dimensionless — consistent.

Note. Version 1 stated $D_{\text{max}} = c \Delta t^* + \sigma \sqrt{\ln(1/\theta)}$; that form is superseded (Version note). The corrected numbers below shift by a few per cent, and the qualitative conclusions survive.

6.1 Numerical Examples (NYSE–Chicago HFT Corridor)

Illustrative parameters (order-of-magnitude, not vendor-specific):

- Required success probability $\theta = 0.8$ (a dependable strategy), giving $\Phi^{-1}(0.2) = -0.842$
- Jitter standard deviation: $\sigma_{\tau, \text{fibre}} = 100 \mu\text{s}$, $\sigma_{\tau, \text{mw}} = 50 \mu\text{s}$
- Effective propagation speeds: $c_{\text{fibre}} \approx 0.67c \approx 200,861 \text{ km/s}$; $c_{\text{mw}} \approx 0.99c \approx 296,795 \text{ km/s}$
- Smoothing lengths: $\sigma_{\text{fibre}} \approx 20.1 \text{ km}$, $\sigma_{\text{mw}} \approx 14.8 \text{ km}$
- Corridor distance (Chicago ↔ NY/NJ DCs): $D \approx 1,180 \text{ km}$ (microwave-optimised path, illustrative)

A. What D_{max} do we get for short-lived signals (at 80% reliability)?

- Fibre, $\Delta t^* = 2 \text{ ms}$: **$D_{\text{max}} \approx 384.8 \text{ km}$** (v1: 411.2; insufficient for 1,180 km)
- Microwave, $\Delta t^* = 2 \text{ ms}$: **$D_{\text{max}} \approx 581.1 \text{ km}$** (v1: 600.6; still short)
- Microwave, $\Delta t^* = 4 \text{ ms}$: **$D_{\text{max}} \approx 1,174.7 \text{ km}$** (v1: 1,194.2). The corridor sits almost exactly on the 4 ms reliability frontier: a dependable microwave strategy *just misses* 1,180 km, while a speculative one ($\theta = 0.2$) reaches the vacuum cap of 1,199 km.

B. What Δt^* is required to trade the corridor dependably? Inverting, $\Delta t^* \geq (D - c_{\text{eff}} \sigma_\tau \Phi^{-1}(1 - \theta)) / c_{\text{eff}} = (D + 0.842 \cdot \sigma) / c_{\text{eff}}$ at $\theta = 0.8$:

- Fibre: $\Delta t \gtrsim 5.96 \text{ ms}^*$ ($v1$: 5.83)
- Microwave: $\Delta t \gtrsim 4.02 \text{ ms}^*$ ($v1$: 3.95)

Interpretation. Two levers expand dependable causal reach: raising **c_eff** (fibre → microwave) moves the mean frontier, and cutting σ_τ shrinks the reliability penalty. These figures are consistent with the measured convergence of Chicago–NY transmission times toward the light-speed bound (Laughlin, Aguirre and Grundfest 2014), and they add a second, less-appreciated implication: *jitter control is worth distance*. At $\theta = 0.8$, every 10 μs of jitter costs a microwave strategy about 2.5 km of reach.

For Earth–Mars (mean one-way delay $\sim 1,200 \text{ s}$), $c \Delta t^*$ -scale distances dwarf any plausible σ_τ contribution: with $\sigma_\tau = 5 \text{ ms}$, the jitter term is of order kilometres against a mean frontier of order 10^8 km .

7 The Delayed-Hedging Cost of Communication Blackouts

Continuous hedging presumes the hedger can observe and rebalance continuously. Across a communication link, rebalancing is impossible during the round-trip blackout $\tau_{rt} = 2d/c$; the writer carries unhedgeable gamma exposure between adjustments. Standard discrete-hedging analysis (Boyle and Emanuel 1980) gives a per-interval P&L standard deviation of approximately $\frac{1}{\sqrt{2}} |\Gamma| S^2 \sigma_A^2 \tau_{rt}$ (with σ_A the asset's return volatility), accumulating over an option's life T , rebalancing each blackout, to

$$\text{Std}_{\text{unhedged}} \approx \sqrt{\frac{1}{2} \tau_{rt} T} |\Gamma| S^2 \sigma_A^2.$$

A CARA/quadratic writer's reservation price exceeds Black–Scholes by a term proportional to the corresponding *variance* (approximately half the risk-aversion coefficient times it). This is a cost of risk bearing, not a violation of no-arbitrage: in a frictionless market, discrete hedging leaves the arbitrage-free price unchanged at leading order while adding variance.

Worked example (Earth-hedged Mars book). ATM call, $S = 100$, $\sigma_A = 30\%$ p.a., $r = 0$, blackout $\tau_{rt} = 2,400 \text{ s}$ (20-minute one-way delay):

Maturity	BS price	Unhedgeable P&L std	As % of premium
1 hour	0.128	0.074	57.7%
1 day	0.626	0.074	11.8%
1 week	1.657	0.074	4.5%
1 month	3.430	0.074	2.2%
1 year	11.924	0.073	0.6%

The delay cost is **maturity-dependent**: decisive for short-dated, high-gamma claims; basis points for long-dated ones. The structural implication connects to the market-design discussion of Section 8.2.3: markets under long delay will not support short-dated continuously-hedged derivatives, and activity migrates to longer tenors and auction formats. A *price-of-risk* premium (expected-return compensation for latency exposure) would additionally require a pricing model with a non-traded latency state variable; it is not derived here and no figure is asserted for it.

8 Empirical and Theoretical Applications

Empirical parallels: Microwave HFT links ($\sim 0.99c$) expand causal overlap compared to fibre ($\sim 0.67c$). Cross-border markets exhibit persistent spreads from latency.

Interplanetary trading (3–22 light-minutes) lies far beyond any terrestrial $D_{\max}$.

Relativistic Market Efficiency Index: $E_c = (\text{causally connected share of market activity}) / (\text{total market activity})$, computed over venue pairs (i,j) weighted by traded volume, a pair counting as connected for a strategy class with persistence Δt^* when $d_{ij} \leq D_{\max}(\Delta t^*)$. Increasing E_c demands energy and infrastructure, revealing a thermodynamic cost of synchrony.

A further testable implication follows from the delayed-observation structure: cross-venue return correlations must vanish at sampling horizons short relative to inter-venue light-time and recover as the horizon lengthens. This is the **Epps effect** (Epps 1979), with a ramp horizon set by physical separation and effective latency rather than by asynchronous ticks alone.

8.1 Illustration: dislocation lifetimes and causal reach

Version 1 of this paper presented the 1 August 2012 Knight Capital incident (SEC 2013) as an empirical validation of $D_{\max}$, arguing that traders beyond ~410 km were causally excluded from the mispricings. That reading is withdrawn: the Knight dislocations unfolded over ~45 minutes, so information crossed the continent thousands of times during the event; distant participation was rationed by venue microstructure (queue position, co-location, risk limits) and by the speed of *decision-making*, not by light cones. What Knight actually illustrates is the cost of autonomous execution outrunning oversight, a governance-latency problem related to, but distinct from, the causal bounds developed here.

The proper illustrations of $D_{\max}$ are dislocations whose lifetimes are comparable to inter-venue light-times: sub-second flash events of the type catalogued by market-data analysts (Nanex 2012) and analysed in the flash-crash literature (Kirilenko, Kyle, Samadi and Tuzun 2017; Menkveld and Yueshen 2019), where Δt^* is measured in single milliseconds. For $\Delta t^* = 2$ ms and fibre parameters as in Section 6.1, dependable causal reach is under 400 km: for such events, venues a continent apart are informationally separate markets, and the geographic distribution of who can respond is set by physics before it is set by skill. Testing this prediction against catalogued sub-second events (response participation as a function of venue distance) is an open empirical exercise, and a sharper one than any single-event anecdote.

8.2 Policy and Market Design Implications

The RAP framework provides a physics-based lens on longstanding debates about market fairness and high-frequency trading; the regulatory implications of light-speed limits were first mapped by Angel (2014), whose analysis of simultaneity ambiguity in "best price" determinations this section builds on. Unlike informational advantages from research or analysis, latency advantages arise purely from spatial position within causal domains, creating a form of structural inequality that is quantifiable through $D_{\max}$. This section explores how regulatory frameworks and market microstructure might adapt to explicitly account for finite information velocity. We emphasise that RAP is a positive framework (describing what *is*), not normative (prescribing what *should be*). The implications discussed below represent logical consequences of causal bounds, leaving policy judgements to regulators and market participants.

8.2.1 Regulatory Considerations Under Causal Constraints

Traditional securities regulation distinguishes public from private information. RAP makes explicit a third category, *causally inaccessible public information*: data that is publicly disclosed but unreachable due to propagation delay. This complicates insider trading doctrine and disclosure requirements. (Angel 2014 identifies the cognate problem for trade sequencing: observers at different locations may legitimately disagree about which of two spacelike-separated events occurred first.)

Speed limits versus access democratisation: Regulatory responses to latency advantages fall into three categories: (a) *Artificial delays* (e.g., IEX "speed bump") homogenise arrival times; (b) *Universal fast access* (subsidised infrastructure) reduces σ_τ and increases c_{eff} ; (c) *Causal transparency* (latency disclosure requirements) enables informed consent.

RAP suggests that option (c), mandating exchanges disclose σ_τ , c_{eff} , and $D_{\max}(\Delta t^*)$, provides a foundation for the others. Traders can assess their causal disadvantage and price it into participation decisions. This parallels nutrition

labelling: information enables choice without restricting market structure.

Cross-border market integration: Two exchanges may trade identical assets yet remain causally fragmented. Consider London–New York equity markets with $\tau \approx 28$ ms one-way on fibre routes (the vacuum floor is 18.6 ms). For price movements with $\Delta t^* < \tau$, these markets are *informationally separate* despite appearing integrated. Regulatory definitions of "consolidated market" implicitly assume $c \rightarrow \infty$ (cf. SEC 2010); finite c requires causal overlap metrics.

The effective integration is not binary but continuous:

$$\text{Integration}_{\text{causal}} = \exp(-\tau^2/2\sigma_{\text{market}}^2),$$

where σ_{market} represents the characteristic time scale of price discovery.

Order routing and best execution: Current best-execution rules focus on price priority. Under RAP, execution quality depends on causal accessibility. A better price at a distant venue may be illusory if $\tau > \Delta t^*$.

Best execution should incorporate latency: traders beyond D_{max} cannot "reach" certain quotes before they update. Execution quality standards should therefore be *conditional on trader location*, much as telecommunications quality-of-service varies geographically.

Off-exchange trading and causal fragmentation: The proliferation of dark pools and alternative trading systems (Zhu 2014; Comerton-Forde and Putniņš 2015) introduces additional causal complexity. Dark pools physically located near public exchanges ($d \approx 100$ m) share similar causal domains as lit markets, but those geographically dispersed create isolated causal islands.

Consider two dark pools separated by $d = 2,000$ km: even if both trade identical assets, they are causally disconnected for sub-10 ms price movements ($\tau = 2000/(0.67 \times 300,000) \approx 10$ ms).

This raises regulatory questions: (a) Should consolidated tape requirements account for causal accessibility, given that reporting delays from distant venues may render price information stale before dissemination? (b) Does dark pool proliferation systematically reduce E_c by fragmenting liquidity into causally isolated venues? (c) For cross-venue arbitrage, should trade-through rules consider whether a "better" price at a distant venue is causally reachable before it updates?

RAP suggests that venue fragmentation has both *spatial and temporal* dimensions: two venues separated beyond D_{max} function as effectively separate markets for high-frequency strategies, even if regulatory frameworks treat them as a unified national market system. Effective integration requires not just electronic connectivity but causal overlap.

Cross-border regulatory arbitrage under causal constraints: International markets face asymmetric regulatory enforcement due to causal structure. A trader in London ($\tau_{\text{NY}} \approx 28$ ms, $\tau_{\text{HK}} \approx 80$ ms) operates within the causal domain of New York but largely outside Hong Kong's for millisecond-scale events. This creates enforcement challenges: manipulative strategies executed from locations with advantageous causal positioning (low τ to multiple jurisdictions) can exploit regulatory gaps before distant authorities detect violations.

Regulatory harmonisation efforts (IOSCO principles, MiFID II/Reg NMS coordination) implicitly assume information symmetry among regulators. RAP reveals this assumption breaks for high-frequency violations: by the time evidence propagates to distant regulators ($\tau \approx$ tens of milliseconds), manipulative positions may already be unwound.

Regulatory coordination should therefore account explicitly for detection delays: (a) shared surveillance infrastructure positioned at causal hubs to minimise $\max\{\tau(\text{hub}, \text{jurisdiction}_i)\}$; (b) "regulatory batch windows" where violations are assessed over intervals $\geq D_{\text{max}}/c_{\text{eff}}$ to ensure all jurisdictions access equivalent information; (c) coordinated circuit breakers with causal propagation delays explicitly modelled.

The RAP framework implies that effective global market regulation requires not just legal harmonisation but *causal synchronisation*: regulatory data and enforcement capabilities must respect the physical topology of information flow.

Enforcement mechanisms and observability: Implementing causal fairness requires monitoring capabilities that respect $D_{\max}$ constraints. Consider latency disclosure requirements: exchanges must report σ_{τ} , c_{eff} , and $D_{\max}(\Delta t^*)$. Verification requires regulators to independently measure network latency, but regulator infrastructure itself faces causal limits. A regulator in Washington, DC measuring latency to NYSE Mahwah ($d \approx 360$ km, $\tau \approx 1.8$ ms on fibre) cannot directly observe sub-millisecond advantages enjoyed by co-located traders ($\tau < 0.5 \mu\text{s}$ at $d \approx 100$ m). This measurement asymmetry suggests two approaches:

(a) *Distributed regulatory nodes* positioned within critical $D_{\max}$ radii (e.g., co-located monitoring at major exchanges), trading transparency for enforceability; (b) *Cryptographic timestamps* with Byzantine-fault-tolerant consensus across geographically distributed validators, providing tamper-proof latency certification without requiring regulators to operate within every causal domain.

Beyond that, "best execution" enforcement under causal constraints requires auditing not only order routing decisions but whether routed venues were causally accessible given market conditions at decision time. This is computationally intensive: for each order, calculate whether alternate venues with better prices satisfied $d \leq D_{\max}(\Delta t^*_{\text{actual}})$, where $\Delta t^*_{\text{actual}}$ is the observed price-quote lifespan. Current surveillance systems lack this causal analysis capability.

These enforcement challenges suggest that regulatory adaptation to RAP principles requires not merely policy changes but substantial infrastructure investment in distributed monitoring, high-resolution timing, and causal topology analysis systems.

8.2.2 Fairness and the Topology of Information Access

RAP reveals that market participants face two distinct forms of information asymmetry: epistemic (skill-based knowledge advantages) and *causal* (physics-based accessibility limits). The latter creates systematic advantages independent of trader sophistication.

Consider two traders separated by distance d attempting to arbitrage a price discrepancy with decay time Δt^* . The distant trader can dependably participate only if $d \leq D_{\max} = c_{\text{eff}} \Delta t^* + c_{\text{eff}} \sigma_{\tau} \Phi^{-1}(1 - \theta)$. Beyond this threshold, arbitrage opportunities vanish before information arrival, creating a "fairness radius" around liquidity venues.

Quantifying spatial privilege: High-frequency trading firms pay substantial fees for exchange co-location, effectively purchasing reduced $\tau(x, x')$. A trader co-located at distance $d_1 \approx 100$ m versus a distant trader at $d_2 = 1,000$ km experiences a latency differential:

$$\Delta \tau = (d_2 - d_1)/c_{\text{eff}} \approx 1,000 \text{ km} / (0.67 \times 3 \times 10^5 \text{ km/s}) \approx 5 \text{ ms}.$$

For strategies with $\Delta t^* \approx 1$ ms, the distant trader is causally excluded. The co-location fee thus represents rent on proximity within $D_{\max}$, analogous to spectrum licensing in telecommunications.

The Relativistic Market Efficiency Index E_c operationalises fairness as causal completeness. Markets with $E_c \approx 1$ approach global information accessibility; those with $E_c \ll 1$ exhibit structural fragmentation. Unlike traditional efficiency metrics that assume instant information, E_c captures the *physical geometry* of information access.

Geographic inequality emerges naturally: Financial centres (New York, London, Hong Kong) become information hubs not only due to liquidity concentration but also due to causal centrality. Peripheral markets face double disadvantage: lower liquidity *and* higher average τ to major venues. This suggests market concentration is partially a consequence of informational topology, not merely network effects or regulation.

8.2.3 Market Microstructure Design in Finite-Speed Regimes

Market design traditionally optimises price discovery and liquidity provision assuming instant information. RAP suggests alternative architectures explicitly designed for causal structure.

Batch auctions with causal windows: Rather than continuous matching, exchanges could batch orders over intervals $T_{\text{batch}} = D_{\text{max}}/c_{\text{eff}}$, treating all orders within the same causal horizon simultaneously. For an exchange targeting continental coverage with $D_{\text{max}} \approx 3,000$ km, $T_{\text{batch}} \approx 15$ ms on fibre. Orders from Boston ($\tau \approx 1.5$ ms) and Chicago ($\tau \approx 5\text{--}6$ ms) arrive within the batch window and compete equally. This eliminates pure proximity advantages while preserving sub-batch latency competition.

Budish, Cramton, and Shim (2015) advocate frequent batch auctions to curb latency arms races, and Petrov (2025) independently proposes batch auctions synchronised within local light cones. RAP provides a physical basis for optimal batch size: T_{batch} should equal the maximum acceptable D_{max} divided by effective propagation speed. Shorter intervals reintroduce causal advantages; longer intervals sacrifice price responsiveness.

Distributed matching engines: Centralised order books implicitly favour traders near the matching engine. An alternative: geographically distributed nodes, each with local matching, synchronised at intervals $\tau_{\text{sync}} > \max\{\tau(\text{node}_i, \text{node}_j)\}$. Orders route to the causally nearest node. This increases E_c by reducing average τ across all traders. The challenge lies in maintaining cross-node consistency, a distributed-systems problem analogous to consensus protocols in blockchain networks.

Dynamic tick sizes: Minimum price increments (ticks) should scale with causal uncertainty. In highly fragmented markets (low E_c), stochastic latency σ_τ creates price noise. Setting $\text{tick}_{\text{min}} \propto \sigma_{\text{price}}(1 + \sigma/c_{\text{eff}} \Delta t^*)$ prevents spurious arbitrage from network jitter while allowing genuine price discovery. This connects microstructure to physical infrastructure: improving network reliability (reducing σ_τ) justifies tighter spreads.

Interplanetary markets (forward-looking): Earth–Mars communication exhibits $\tau \approx 3\text{--}22$ minutes, far exceeding typical Δt^* . Continuous markets are impossible; instead, periodic batch auctions synchronised to light-time become necessary, and the maturity structure of Section 7 implies the derivative book migrates to long tenors. Orders submitted on Earth lock during transit; Mars market prices at auction close reflect Earth information from τ minutes prior. This represents an extreme regime where RAP is not an approximation but a fundamental constraint. As human activity extends beyond Earth, financial infrastructure must be designed with relativistic principles from inception.

8.2.4 Batch Auctions as Settlement-Frame Enforcement

Batch auctions solve a second problem beyond causal fairness. Carvalho and Gaspar (2021) show that participants in different relativistic reference frames accumulate different proper times, so a common **settlement reference frame** is needed to prevent valuation arbitrage between clocks. Batch windows supply one: define T_{batch} in the exchange's reference frame and mandate settlement-frame timestamps on all orders, and every participant transacts in a single time coordinate regardless of their local clock.

At today's velocities the requirement is a convention rather than a correction. Terrestrial $v/c \approx 10^{-8}$ makes dilation negligible ($\gamma \approx 1 + 10^{-16}$), and where dilation does bind — GPS satellites, at ~ 38 $\mu\text{s/day}$ — engineering already absorbs it. The batch window need only dominate both constraints: $\Delta t_{\text{batch}} \geq \max\{D_{\text{max}}/c_{\text{eff}}, \text{clock-synchronisation tolerance}\}$. The institutional layer this presupposes is now being built: the Coordinated Lunar Time (LTC) programme, directed by the US OSTP in 2024 with a NIST technical framework, is constructing the relativistically corrected reference timescale to which off-Earth settlement frames can anchor (OSTP 2024).

The batch auction is therefore doing double duty: a causal fairness mechanism (proximity inside D_{max} confers no advantage within a window) and a temporal harmonisation mechanism (all participants price in one time coordinate). Both roles support frequent batch auctions as core infrastructure for physics-constrained markets.

8.2.5 Open Questions for Policy and Design

RAP formalises but does not resolve several normative questions:

1. Should causal domain access near exchanges be treated as a public resource requiring fair allocation, similar to electromagnetic spectrum?
2. Is infrastructure investment to reduce τ socially beneficial (improving E_c) or privately rent-seeking (advantaging a few)?
3. Should co-location fees be uniform or distance-adjusted to neutralise proximity advantages?
4. Can market makers justify priority access if their liquidity provision genuinely improves E_c for all participants?

These questions require empirical investigation of how latency advantages translate to price discovery versus pure rent extraction. RAP provides the framework; data will inform policy.

Note: Full treatment of optimal mechanism design under causal constraints is deferred to future work. The above discussion outlines implications; formal policy modelling requires coupling RAP with game-theoretic market design, which exceeds the scope of this paper.

9 Future Research Pathway

Empirical tests: natural experiments via submarine-cable latency shifts (documented reroutings produce tens-of-milliseconds jumps usable as exogenous shocks); the Epps-ramp prediction of Section 8 against cross-venue tick data; asynchronous Monte Carlo simulations on causal lattices.

Institutional and policy material will appear separately. Subsequent work applies the causal framework to the financing of assets in transit under communication delay (the regime where collateral can be neither observed nor repossessed within the cone), which the author regards as the nearest commercially relevant application of these constraints.

Appendix A: Telegraph Micro-Foundation for Finite-Velocity Information Transport

Persistent random walks provide the micro-foundation for finite-velocity information transport: why market information moves at c_{eff} with damping rather than diffusing instantaneously. The one- and three-dimensional cases carry different coefficients and should not be mixed:

(i) One dimension (Kac). A particle moves at speed c and *reverses* direction at Poisson rate a . The density satisfies the telegraph equation

$$\partial_{tt}\rho + 2a \partial_t\rho = c^2 \partial_{xx}\rho,$$

so matching $\partial_{tt}\delta + 2\gamma\partial_t\delta = c^2\Delta\delta$ gives $\gamma = a$ and propagation speed c , with diffusive limit $v = c^2/(2a)$ (Kac 1974; Goldstein 1951).

(ii) Three dimensions (isotropic redirection). A particle moves at speed c and is *re-randomised isotropically* at rate λ . From the kinetic equation $\partial_t p + v \cdot \nabla p = -\lambda p + \lambda \int p dv'/4\pi$, the current relaxes at rate λ (not 2λ):

$$\partial_t \rho + \nabla \cdot j = 0, \quad \partial_t j + \lambda j = -\frac{c^2}{3} \nabla \rho,$$

yielding

$$\partial_{tt}\rho + \lambda \partial_t\rho = \frac{c^2}{3} \Delta\rho.$$

Matching the damped-wave form gives $\gamma = \lambda/2$ and **effective propagation speed** $c/\sqrt{3}$, with diffusive limit $v = c^2/(3\lambda)$ (Green–Kubo: $v = (1/3)\int \langle v(0) \cdot v(t) \rangle dt = c^2/(3\lambda)$). For traded-price applications of telegraph dynamics with rigorous equivalent martingale measures, see Ratanov (2007) and Kolesnik & Ratanov (2013).

Monte Carlo verification. 40,000 walkers in 3D with $\lambda = 1$, $c = 1$, $T = 40$: measured long-time diffusivity $v = 0.322$, against $c^2/(3\lambda) = 0.333$ (the residual is the expected ballistic correction) and against the superseded note's claimed $c^2/(6\lambda) = 0.167$, which is rejected. With persistence time $\tau_p = 1/\lambda \approx 60$ s, the identification $\gamma = \lambda/2$ gives $\gamma \approx 1/120 \text{ s}^{-1} \approx 0.0083 \text{ s}^{-1}$, the value used numerically throughout.

Calibration status. An autocorrelation fit previously described as "empirical calibration using DSN data for Mars" was in fact a self-consistency check on synthetic data generated from the assumed model, and is labelled as such in Appendix B. A genuine calibration would fit the autocorrelation of *residual* latencies (observed minus deterministic geometric delay) from published Deep Space Network ranging or telemetry logs; until that is done, all values of γ and σ_τ in this paper are illustrative parameters.

Appendix B: Reproducibility

Every numerical result printed in this paper is reproducible from the two functions below: the corridor figures of Section 6.1 from `D_max`, and the hedging table of Section 7 from `delayed_hedging_cost`. The damped-wave information-field simulator (superseded as a price model by Proposition 1, retained as an information-transport tool) and its Newmark- β stability analysis are provided in the SSRN ancillary notebook rather than here, since no result in this paper depends on them.

```
def D_max(c_vac, c_eff, Delta_t_star, sigma_tau, theta):
    """Corrected reach bound (Section 6); theta = required survival probability."""
    from scipy.stats import norm
    return min(c_vac * Delta_t_star,
               c_eff * Delta_t_star + c_eff * sigma_tau * norm.ppf(1 - theta))

def delayed_hedging_cost(S, sigma, T, tau_rt, r=0.0):
    """Unhedgable P&L std for an ATM call rebalanced each blackout tau_rt (years). Section 7."""
    from math import log, sqrt, exp, pi, erf
    N = lambda z: 0.5*(1+erf(z/sqrt(2)))
    phi = lambda z: exp(-z*z/2)/sqrt(2*pi)
    d1 = (r + 0.5*sigma**2)*T / (sigma*sqrt(T))
    C = S*(N(d1) - N(d1 - sigma*sqrt(T)))
    G = phi(d1)/(S*sigma*sqrt(T))
    std_total = sqrt(0.5*tau_rt*T) * G * S**2 * sigma**2
    return C, std_total, std_total/C
```

```
# SYNTHETIC SELF-CONSISTENCY CHECK (not an empirical calibration).
# Data are generated from acf_model with gamma = 1/60 embedded; the fit
# recovers the input. A real calibration would fit the autocorrelation of
# residual latencies from published DSN ranging/telemetry logs.
import numpy as np
from scipy.optimize import curve_fit

def acf_model(t, gamma):
    return np.exp(-2 * gamma * t)
```

```

np.random.seed(42)
t_data = np.linspace(0, 1000, 200)
acf_data = np.clip(acf_model(t_data, gamma=1/60)
                  + np.random.normal(0, 0.05, len(t_data)), 0, 1)
(popt, pcov) = curve_fit(acf_model, t_data, acf_data, p0=[0.01])
print(f"recovered gamma = {popt[0]:.4f} 1/s (input was {1/60:.4f}) -- self-consistency only")

```

No simulated magnitude in this paper is presented as an empirical or market quantity; the synthetic block above is retained in the manuscript as the documentary record of the calibration mislabelling corrected in this version (see the Version note).

Appendix C: Glossary of Key Terms

This glossary provides accessible definitions of technical concepts for readers unfamiliar with stochastic calculus, measure theory, or relativistic physics.

Mathematical & Financial Terms

Arbitrage The simultaneous purchase and sale of an asset to profit from price differences in different markets or at different times. In efficient markets, arbitrage opportunities should be quickly eliminated. RAP demonstrates that geographic separation can sustain price differentials beyond the reach of distant traders.

Brownian Motion (Wiener Process) A continuous random process where the position at time t is the sum of many small random steps. Mathematically: increments are normally distributed, independent, and satisfy $W_0 = 0$. In finance, models the unpredictable "random walk" of asset prices.

Causal Domain The region of spacetime from which information can reach a given point. Analogous to the "past light cone" in physics. In markets, a trader's causal domain contains all price information accessible given communication delays.

Causal Filtration ($\mathcal{F}_{\Lambda(t,x)}$) A mathematical structure (σ -algebra) representing all information accessible to an observer at position x and time t , accounting for finite propagation speed. Formally: the collection of all events within the past light cone $\Lambda(t,x)$. Classical finance assumes the full filtration (all information available everywhere).

Delayed-Observation Chain Rule Classical Itô calculus applied to a time-shifted process: when the locally observed price is the light-delayed remote price $S_{t-\tau}(x)$, the differential of $f(S^{\text{obs}}, t)$ follows by ordinary Itô with arguments shifted by τ . Version 1 called this a "relativistic Itô lemma"; the present name reflects that it is a change of variables within standard theory (see Mohammed 2013 for the rigorous treatment of delayed stochastic systems).

Diffusion A process where quantities spread out over time, typically following a parabolic partial differential equation (heat equation). Classical asset price models use diffusion with instantaneous propagation; finite-velocity alternatives exist (telegraph processes, Kolesnik & Ratanov 2013).

Discounted Price The current price divided by the value of a risk-free bond that pays \$1 at maturity: S_t/B_t . The Fundamental Theorem of Asset Pricing states that discounted prices are martingales under a risk-neutral measure.

Equivalent Measure Two probability measures are equivalent if they assign zero probability to exactly the same events (though probabilities may differ for non-zero events).

Filtration A growing collection of information over time. Formally: an increasing sequence of σ -algebras. Think of it as an "information diary": $\mathcal{F}_t$ contains all information known up to time t . RAP uses causal filtrations where information growth

is limited by propagation speed.

Fundamental Theorem of Asset Pricing (FTAP) States that a market is arbitrage-free if and only if there exists a risk-neutral probability measure under which discounted asset prices are martingales. Under delayed and restricted information the theorem has been established by Kabanov and Stricker (2006) and successors; RAP specialises that setting to light-cone geometry.

Itô Lemma The "chain rule" for stochastic calculus. Allows computation of how a function of a stochastic process evolves. For $f(S_t, t)$, the Itô lemma includes a second-order correction term: $\partial^2 f / \partial S^2 \times (\text{volatility})^2$.

Light Cone In physics: the surface in spacetime representing all events reachable by a signal travelling at light speed from a given point. In RAP: analogously defines the boundary of causally accessible information. Events outside the light cone are informationally unreachable.

Martingale A stochastic process where the expected future value, given all past information, equals the current value: $E[S_T | \mathcal{F}_t] = S_t$. Represents a "fair game" with no predictable drift.

Measure-Theoretic Relating to the mathematical framework for defining probabilities rigorously. Measure theory provides foundations for probability spaces, random variables, and expectations.

Predictable Strategy A trading strategy where portfolio adjustments at time t depend only on information available strictly before t . In RAP, strategies must be predictable with respect to causal filtrations.

Quadratic Variation For a stochastic process, the sum of squared increments over infinitesimal intervals. For Brownian motion: $[W, W]_t = t$. Appears in the second-order term of Itô's lemma. Processes with differentiable paths have zero quadratic variation, a distinction that matters for which dynamics can support martingale prices (see Section 4).

Risk-Neutral Measure (Q) A probability measure under which all assets, when discounted, grow at the risk-free rate. Used for pricing derivatives without knowing investors' risk preferences.

Self-Financing Strategy A portfolio where changes in value arise only from price movements and reinvested gains/losses, not from external cash injections.

Sigma-Algebra (σ -algebra) A collection of events (subsets) closed under countable unions, intersections, and complements. Represents "measurable information".

Square-Integrable A function or process with finite expected squared value: $E[X^2] < \infty$.

Stochastic Process A family of random variables indexed by time: $\{S_t : t \geq 0\}$.

Strategy (Trading Strategy) A rule specifying portfolio holdings over time. Admissible strategies in RAP respect causal constraints.

Physics & Relativistic Terms

Causal Structure The network of which events can influence which other events, determined by maximum signal speed. In physics: the speed of light. In RAP: the effective information speed c_{eff} , hard-capped by c .

Hyperbolic PDE Partial differential equations of wave type, with finite propagation speed. Standard form: $\partial^2 u / \partial t^2 = c^2 \nabla^2 u$. Contrast with parabolic PDEs (diffusion) that have infinite propagation speed. In RAP these model information transport; Section 4 explains why they cannot model the traded price directly.

Information Geometry Using geometric concepts (manifolds, metrics, geodesics) to study information structures. RAP treats the market as a 4-dimensional manifold with a Minkowski-like metric encoding information causality.

Light Cone (Past/Future) The set of all points in spacetime that can send signals to (past cone) or receive signals from (future cone) a given event, assuming signals travel at speed c .

Minkowski Metric The metric of special relativity: $ds^2 = c^2 dt^2 - dx^2 - dy^2 - dz^2$. Separates timelike (causal connection possible) from spacelike (causally disconnected) intervals. RAP uses this as an informational metric.

Propagation Speed The rate at which information travels through a medium. In RAP: effective information speed c_{eff} depends on communication infrastructure (fibre: $\sim 0.67c$, microwave: $\approx c$ in air), always bounded by vacuum c .

Retarded Green's Function The solution to a wave equation representing the effect of a source at earlier time t' on a field at later time t . Has support only on the past light cone.

Spacelike Separation Two events are spacelike-separated if the spatial distance exceeds what light could travel in the time difference: $|x_1 - x_2| > c|t_1 - t_2|$. Such events cannot causally influence each other. In RAP: traders at spacelike separation cannot arbitrage against each other.

Spacetime The four-dimensional continuum combining three spatial dimensions and one time dimension.

Telegraph Equation A hyperbolic PDE: $\partial^2 u / \partial t^2 + (1/\tau) \partial u / \partial t = c^2 \nabla^2 u$. Describes finite-speed diffusion with damping. In finance, telegraph processes underpin a developed option-pricing literature (Kolesnik & Ratanov 2013; Ratanov 2007); Trzetrzelewski (2017) derives a relativistic Black–Scholes model from the same equation. See Appendix A for the corrected coefficients.

Timelike Separation Two events are timelike-separated if a signal travelling at speed $\leq c$ can connect them. Such events can influence each other.

Wave Equation The prototypical hyperbolic PDE: $\partial^2 u / \partial t^2 = c^2 \nabla^2 u$. Solutions propagate at speed c with sharp wavefronts.

RAP-Specific Terms

Causal Arbitrage Arbitrage constrained by finite information propagation. Two traders can dependably arbitrage only if separated by $d \leq D_{max} = \min\{c\Delta t^*, c_{eff} \Delta t^* + c_{eff} \sigma_\tau \Phi^{-1}(1-\theta)\}$. Beyond this distance, price differentials decay before information arrives with the required reliability.

Causal Overlap The intersection of causal domains of two traders. If overlap is empty (spacelike separation), no shared information exists for synchronised trading.

Causal Smoothing Length (σ) The standard deviation of network latency jitter, in distance units ($\sigma = c_{eff} \times \sigma_\tau$). Represents stochastic dispersion of signal arrival around the mean transit time. Typical HFT values: $\sigma \approx 2\text{--}30$ km.

Causality-Bounded Pricing (restated) The no-arbitrage statement of Section 3: sufficiency of an equivalent martingale measure transfers to causal filtrations by conditioning; the converse is the delayed-information FTAP (Kabanov & Stricker 2006), specialised to light-cone geometry. (Version 1's name "CBPT", presented as a new theorem, is retired.)

D_{max} (Maximum Causal Arbitrage Distance) The furthest distance over which arbitrage meets a required success probability θ : $D_{max} = \min\{c\Delta t^*, c_{eff}\Delta t^* + c_{eff}\sigma_\tau\Phi^{-1}(1-\theta)\}$. Reliability ($\theta > 1/2$) shrinks reach below the mean frontier; speculation ($\theta < 1/2$) extends it toward, never beyond, the vacuum cone.

Delayed-Hedging Cost The unhedgeable P&L variance accumulated during communication blackouts (Section 7); maturity-dependent, decisive when option maturity is comparable to the blackout.

Obstruction (Proposition 1) Hyperbolic (second-order-in-time) log-price dynamics produce finite-variation prices with no martingale part, so no equivalent martingale measure exists for the discounted price. Causality must enter through delayed filtrations or noise covariance instead (Section 4).

Relativistic Market Efficiency Index (E_c) Volume-weighted share of venue pairs that are causally connected for a given strategy class: $E_c \approx 1$ indicates near-global information access; $E_c \ll 1$ indicates fragmented markets.

Market Microstructure Terms

Batch Auction A trading mechanism that collects orders over a fixed interval and executes them simultaneously at a single clearing price. RAP suggests batch windows $T_{\text{batch}} = D_{\text{max}}/c_{\text{eff}}$ to equalise traders at different distances (cf. Budish et al. 2015; Petrov 2025).

Best Execution Regulatory requirement that brokers obtain the most favourable terms for client orders. RAP perspective: must account for causal accessibility, since distant "better" prices may update before orders arrive (cf. Angel 2014).

Co-location Placing trading servers physically near exchange matching engines to minimise latency. RAP framework: co-location purchases spatial proximity within D_{max} .

Dark Pool Private trading venues where order information is not publicly displayed. Raises RAP questions: do dark pools create causally isolated markets? Does fragmentation reduce E_c ?

Latency The time delay between sending a message and receiving a response. Determined by physical distance and propagation medium.

Latency Arbitrage Exploiting time delays in price updates across different venues. RAP shows this is fundamentally bounded by D_{max} .

Market Maker An entity providing liquidity by continuously offering to buy (bid) and sell (ask). RAP question: does priority access improve E_c for all participants?

Matching Engine The core system at an exchange that matches buy and sell orders. Physical location determines τ for all traders.

Microstructure The study of how trading mechanisms, market design, and institutional details affect price formation. RAP contributes physics-based constraints to microstructure theory.

Tick Size The minimum price increment for trading. RAP suggests: $\text{tick}_{\text{min}} \propto \sigma_{\text{price}}(1 + \sigma/c_{\text{eff}} \Delta t^*)$ to prevent spurious arbitrage from network jitter.

References

- Angel, J. J. (2014). When finance meets physics: The impact of the speed of light on financial markets and their regulation. **Financial Review**, 49(2), 271–281.
- Arriojas, M., Hu, Y., Mohammed, S.-E. A., & Pap, G. (2007). A delayed Black and Scholes formula. **Stochastic Analysis and Applications**, 25(2), 471–492.
- Barndorff-Nielsen, O. E., Benth, F. E., & Veraart, A. E. D. (2018). *Ambit Stochastics*. Springer.
- Bloss, M., & Kleinknecht, M. (2016). Option valuation under the effect of time dilation. **SSRN Working Paper 2716997**.
- Boyle, P. P., & Emanuel, D. (1980). Discretely adjusted option hedges. **Journal of Financial Economics**, 8(3), 259–282.
- Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. **Quarterly Journal of Economics**, 130(4), 1547–1621.
- Carvalho, V. H., & Gaspar, R. M. (2021). Relativistic option pricing. **International Journal of Financial Studies**, 9(2), 32.

- Comerton-Forde, C., & Putniņš, T. J. (2015). Dark trading and price discovery. **Journal of Financial Economics**, 118(1), 70–92.
- Delbaen, F., & Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. **Mathematische Annalen**, 300(1), 463–520.
- Epps, T. W. (1979). Comovements in stock prices in the very short run. **Journal of the American Statistical Association**, 74(366), 291–298.
- Goldstein, S. (1951). On diffusion by discontinuous movements and the telegraph equation. **Quarterly Journal of Mechanics and Applied Mathematics**, 4(2), 129–156.
- Harrison, J. M., & Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. **Journal of Economic Theory**, 20(3), 381–408.
- Haug, E. G. (2004). Space-time finance. **Wilmott Magazine**, July 2004.
- IOSCO (2011). *Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency*. Consultation Report, CR02/11.
- Kabanov, Y., & Stricker, C. (2006). The Dalang–Morton–Willinger theorem under delayed and restricted information. *In Memoriam Paul-André Meyer, Séminaire de Probabilités XXXIX* (M. Émery & M. Yor, eds.), Lecture Notes in Mathematics 1874, Springer, 209–213.
- Kac, M. (1974). A stochastic model related to the telegrapher's equation. **Rocky Mountain Journal of Mathematics**, 4(3), 497–509.
- Kirilenko, A. A., & Lamacie, G. (2015). Latency and asset prices. **SSRN Working Paper 2546567**.
- Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. **Journal of Finance**, 72(3), 967–998.
- Kolesnik, A. D., & Ratanov, N. (2013). *Telegraph Processes and Option Pricing*. Springer.
- Krugman, P. (2010). The theory of interstellar trade. **Economic Inquiry**, 48(4), 1119–1123. (Written 1978.)
- Laughlin, G., Aguirre, A., & Grundfest, J. (2014). Information transmission between financial markets in Chicago and New York. **Financial Review**, 49(2), 283–312.
- Lerner, P. B. (2021). Transmission of trading orders through communication line with relativistic delay. **International Journal of Financial Studies**, 9(1), 12.
- Menkveld, A. J., & Yueshen, B. Z. (2019). The flash crash: A cautionary tale about highly fragmented markets. **Management Science**, 65(10), 4470–4488.
- Mohammed, S. E. A. (2013). *Stochastic Functional Differential Equations*. Springer.
- Nanex (2012). Knight Capital's \$440 Million Error. Market analysis report, August 1, 2012.
- OSTP (2024). Memorandum on celestial time standardisation. Executive Office of the President, Office of Science and Technology Policy, 2 April 2024.
- Petrov, V. (2025). The relativistic-chaotic market hypothesis: On the physical impossibility of perfect informational efficiency. **MPRA Paper 123835**.
- Ratanov, N. (2007). A jump telegraph model for option pricing. **Quantitative Finance**, 7(5), 575–583.
- SEC (2010). *Concept Release on Equity Market Structure*. Release No. 34-61358, File No. S7-02-10.
- SEC (2013). In the Matter of Knight Capital Americas LLC. Administrative Proceeding File No. 3-15570, October 16, 2013.
- Trzetrzelewski, M. (2017). The relativistic Black–Scholes model. **EPL (Europhysics Letters)**, 117(3), 38004.
- Wissner-Gross, A. D., & Freer, C. E. (2010). Relativistic statistical arbitrage. **Physical Review E**, 82(5), 056104.
- Zhu, H. (2014). Do dark pools harm price discovery? **Review of Financial Studies**, 27(3), 747–789.

AUTHOR DECLARATION This paper represents independent research conducted by the author. The views and opinions expressed are solely those of the author and do not represent the views of any organisation, institution, or employer. The author received no external funding for this work. No confidential or proprietary information has been used in its preparation. Conflicts of interest, where any exist, are disclosed in the body of the paper.

AUTHOR DECLARATION ON AI ASSISTANCE The author used artificial-intelligence tools to assist with language refinement, code debugging, mathematical verification, and literature search. All ideas, analysis, and conclusions are the author's own; all AI-assisted text and computations were reviewed and validated by the author before inclusion; and the author accepts full responsibility for the content of this paper.

VERSION NOTE Version 2.0 (July 2026) revises the 1 November 2025 original and **consolidates the companion technical note** ("Relativistic Asset Pricing: Mathematical Proofs and Numerical Implementation", SSRN 5721142, 7 November 2025) into this single paper; that record now carries a supersession notice and erratum. The title drops the "(Part I)" designation: the programme is consolidated to this paper, and no further numbered instalments are planned.

Principal corrections carried by this version. From the original Part I: the no-arbitrage statement is restated with its sufficiency direction proved by conditioning and its converse attributed to the delayed-information FTAP literature (Kabanov & Stricker 2006), replacing Version 1's presentation of a new "CBPT"; the correlation kernel is corrected to carry hard support inside the vacuum light cone; $D_{\max}$ is rederived from signal-survival probability, replacing the superseded form $c\Delta t^* + \sigma\sqrt{\ln(1/\theta)}$, with corrected corridor numerics (Section 6.1) and a sign correction in the reliability regime; the "relativistic Itô lemma" is restated as a delayed-observation chain rule within classical theory; the Knight Capital section is withdrawn as a validation and replaced with a corrected account of what such events do and do not illustrate; the technical lineage (Wissner-Gross & Freer; Angel; Ratanov/Kolesnik; Trzetrzelewski; Kirilenko & Lamacie; Lerner; Petrov; Laughlin et al.) is cited; and E_c is given an operational definition.

From the consolidated technical note: its central construction — a discounted-price martingale on hyperbolic log-price dynamics via Girsanov — is withdrawn and replaced by the obstruction result (Proposition 1, Section 4) explaining why it cannot work; the three-dimensional telegraph coefficients are corrected ($\gamma = \lambda/2$, effective speed $c/\sqrt{3}$, diffusivity $c^2/(3\lambda)$, Monte Carlo-verified; Appendix A); the note's simulated " premia " (0.2% Moon; 7.8–62.5% Mars; 500% neutron star; 1,499.8% Proxima) are withdrawn as artifacts, being linear in an arbitrary volatility input and not premia in any pricing sense, and are replaced by the derived, maturity-dependent delayed-hedging cost of Section 7; an exercise previously described as DSN calibration is relabelled a synthetic self-consistency check (Appendix B); a misattributed reference is corrected (arXiv:1009.3526 is Da Prato, Jentzen & Röckner, *A mild Itô formula for SPDEs*); code is updated for current SciPy and relabelled; and material on Kerr geodesics, proprietary models, and commercial deployment is removed as out of scope. Conclusions that depended on superseded or withdrawn statements have been restated or removed.