

Capacity and Scalability Evidence for Short-Dated BTC Event Markets

Martial Systems LLC · Research methodology note · KXBTC15M-style instruments

Public site version 1.29 · Generated 2026-08-19 05:05 UTC

Experimental research documentation. Not investment advice. Not a live trading track record. Aggregate methods only.

Revisions (2026-08-19)

- **2026-07-31:** Plan locked. Phase 0 L2 pipe. Public panel and this PDF first published.
- **2026-08-01:** Phase 1 VWAP pot ladder. Phase 0 public.
- **2026-08-02:** Phase 2 fill and footprint. Phase 3 Bracketeer. Phase 4 liq-cap Path P&L.
- **2026-08-03:** Diligence pack. Phases 0 to 4 complete on the public site.
- **2026-08-19:** Dated revisions list added. Implementation dates were missing from this PDF.

Abstract. One-contract shadow Path P&L is appropriate for **edge discovery** under fixed entry gates. It is **not** a valid extrapolation of large notional deployment: larger orders walk the book, consume visible depth, and face product-level concentration (typically one active 15-minute strike). This note summarizes a four-phase **capacity evidence layer**—static Level-2 walks, fill/footprint models, multi-window schedules, and Path P&L under explicit size caps—so diligence questions about scale can be answered without inventing linear size ladders or modifying frozen shadow exit policies.

1. Public research snapshot (aggregates)

The metrics below match the public board's high-level aggregates. Unless noted, P&L is on a **1-contract** shadow path with unique last-write windows. These figures measure signal research maturity; they do not, by themselves, establish large-pot capacity.

Metric	Value	Role in evaluation
Unique settled windows	2167	Sample maturity (edge discovery)
Shadow first_hot (markets)	1275	Entry-level exit research unit
win65 unique tickers	231	Freeze bar before exit policy promotion
Path P&L (1-ct cooldown)	\$5.78	Baseline equity series — pre-capacity
Capacity program status	Phase 4: Phases 0–4 complete — static L2 capacity evidence layer	phase 4 complete
L2 snapshots (disk only)	n≈242443 · 10.285 MB · collecting	Phase 0 input; full books not published

2. Scope: edge discovery vs capacity evidence

The research stack separates decision logic from execution, defaults to paper/shadow operation, and publishes aggregate Path P&L with counterfactual exit comparisons. That design supports **signal evaluation under fixed gates**. It is not a capital-allocation study.

The capacity layer is deliberately parallel. It addresses questions such as:

- At pot size \$X, what is fee-and-drag-aware EV after walking the book (entry + exit VWAP)?
- What fraction of visible depth / window volume does recommended size consume (footprint)?
- Does Path P&L stay net-positive under hard liquidity caps a buyer can toggle?
- If one 15-minute book is thin, how do we schedule capital across windows without pretending independent Kelly on highly correlated binaries?

Legacy shadow exits stay frozen. Existing δ / hybrid / hold-to-settle counterfactuals and the public Shadow WR chart are not retuned for this work. New fill and capacity simulations ship as separate modules and logs so historical exit research remains comparable.

3. Phased capacity program

Each phase has a deliverable and a stop for write-up / sanity check. We do not skip the L2 data pipe: without logged depth, VWAP capacity is fiction.

Phase	Name	Buyer proof	Status
0	L2 data pipe	Disk-only Level-2 bid ladders while research runs (top-N levels). Input for all later sims.	Shipped — logging side-channel; collecting
1	Slippage & book penetration	Pot ladder (\$100–\$5k): entry/exit VWAP, slip ¢, penetration %, fee/drag EV. Kills 1-ct × N math.	Shipped
2	Passive vs aggressive + footprint	New fill models (not old shadow). Footprint gate (size ≤ 12% of visible depth as volume proxy).	Shipped
3	Multi-window capital	Horizontal scale: split pot across windows / markets so per-book penetration stays under caps. Product-aware (often one strike per 15m).	Shipped
4	Liquidity-cap Path P&L	Parallel path_pnl_liq_cap_* equity under hard contract/depth/VWAP-slip caps (uncapped public equity untouched).	Complete

Program status. Phases 0–4 are complete on the current research host sample. Subsequent work is sample refresh (re-running the same reports as L2 and windows grow), not additional phase invention, unless a new diligence question is specified.

4. Phase 1 results (VWAP pot ladder)

Sample: **n=242443** L2 snapshots · **787** unique markets · report 2026-08-15T05:56:22 UTC. Primary envelope: ≥80% full fill and median entry penetration ≤15%.

Side	Max pot (liquidity envelope)	Median slip / pen at envelope
YES	\$2,000	0.63¢ slip · 12.9% pen
NO	\$1,000	0.31¢ slip · 6.2% pen

Secondary (scenario EV@p=0.65 still median-positive under fill rules): up to about **\$2,000** — scenario p only, not live ensemble edge. Suggested liquidity pots: YES \$2,000 · NO \$1,000.

Assumptions: static L2 (no refill); beyond top-N levels unfillable; aggressive entry + exit on the same snapshot; fee/drag EV uses a p grid.

5. Phase 2 results (fill models + footprint)

Sample: **n=242443** snapshots · **787** markets · footprint f=12% of visible depth (volume proxy until trade tape). Modes: **capacity_aggressive** (ask walk) vs **capacity_passive** (rest at bid; time×depth fill proxy). Legacy shadow_exits untouched.

Cell	Med fill	Med slip	Footprint OK %
yes 10 aggressive	100%	0.00¢	96.8%
yes 10 passive	100%	0.00¢	97.1%
no 10 aggressive	100%	0.00¢	97.1%
no 10 passive	100%	0.00¢	96.8%

Phase 1 **\$ pot / penetration** remains the tighter capital envelope on many books; Phase 2 confirms small contract sizes (e.g. 10 ctr) usually clear footprint and fill under the static proxies.

6. Phase 3 results (multi-window Bracketeer)

Sample: **n=242443** L2 snapshots · **787** markets · per-window penetration cap **≤5%** of visible ask depth · budget fill **≥80%**. Horizontal scale only (one KXBTC15M strike per 15m window). Correlation haircut applied so multi-window is not sold as independent Kelly.

Side	Median \$/window	~1 hour (4w)	~24 hour (96w)
YES	\$740	\$2,961	\$71,073
NO	\$776	\$3,104	\$74,499

YES: \$1k≈2 windows · \$10k≈14 windows · NO: \$1k≈2 windows · \$10k≈13 windows

Assumptions: static L2; consecutive windows share BTC regime (correlation haircut); not a live multi-window allocator.

7. Phase 4 results (liquidity-cap Path P&L)

Base path = public equity construction (unique last-write **sim_pnl_cooldown** 1ctr). **n=1223** unique windows · 1ctr path P&L **\$18.01**. Parallel **path_pnl_liq_cap_*** series only — uncapped Path P&L and shadow_exits unchanged.

Series	Path \$	Max DD \$	+rate
1	\$18.01	\$2.63	11.4%
10	\$180.06	\$26.26	11.4%
50	\$900.54	\$131.32	11.4%
100	\$1,803.39	\$262.64	11.4%
depth	\$1,572.13	\$262.64	11.4%
vwap_slip_10	\$180.06	\$26.26	11.4%
vwap_slip_50	\$900.54	\$131.32	11.4%
vwap_slip_100	\$1,488.96	\$317.47	11.4%

Survival under forced liquidity constraints: **8/8** series with path P&L > 0 (depth and VWAP-slip correctly cut naive 1ctr×N linear scale). Buyer visual: edge can survive size/depth/slip constraints without claiming live multi-contract fills.

Assumptions: historical path counterfactual; static depth/slip proxies; no dynamic refill or signaling model.

8. Product-law control plane (GraphForge)

Capacity work is fail-closed under a shared GraphForge runner (engine ≥0.2.3). Domain graphs live in the BTC suite; the engine provides **require_law / graphforge-gate** so SPY and BTC share one update path (version pin + git pull, not pip install). Capacity scripts call the gate on start: no shadow_exits edits, no MODEL_VERSION / gate retunes, phase order enforced. This is ops verification — not a separate trading product and not Path P&L.

9. Program status & diligence pack

Shipped (Phases 0 to 4 complete, 2026-08-03)

- Persistent L2 history (Phase 0) — disk-only full books, not public rows.
- Symmetric entry + exit VWAP at multiple pot sizes with fee/drag EV (Phase 1).
- Separate capacity fill models and footprint metrics (Phase 2).
- Multi-window Bracketeer schedule under per-book penetration caps (Phase 3).
- Liquidity-capped Path P&L parallel series **path_pnl_liq_cap_*** (Phase 4).

- GraphForge product-law gates on capacity scripts (shadow frozen).

Diligence pack checklist

- L2-backed VWAP pot ladder with fee-and-drag-aware EV by pot size.
- Footprint % distribution at recommended operational size.
- Multi-window schedule if single-book capacity binds.
- Path P&L under explicit liquidity caps (contract / depth / VWAP-slip).
- Explicit limits: static book walk $\neq$ full dynamic market impact; signaling and adverse selection are labeled assumptions until live size experiments.

Still true of production research (edge discovery)

- Size ladders (1/5/10/25 contracts and risk% bankrolls) with top-of-ask depth caps — still largely linear PnL $\times$ contracts for discovery; capacity evidence is the separate diligence layer above.
- Public Path P&L and shadow exit comparison remain 1-contract aggregate-only (shadow frozen; not replaced by liq-cap series).

10. Honest limits

Public pages remain **aggregate-only** (no tickers, strikes, settle spots, entry prices, or trade tapes). Full L2 streams stay on research hosts and are not published. Static order-book penetration is a necessary but incomplete model of liquidity; refill, queue priority, and competitor response require additional live or A/B evidence. This document describes a research-to-product evaluation path — not a solicitation, performance guarantee, or claim of exchange affiliation.

Related public materials: pre-correction baseline PDF (realized-vol method note); live board at the BTC 15m research page. Contact: martialsys@gmail.com · Martial Systems LLC.

Copyright (c) 2026 Martial Systems LLC. All rights reserved. Proprietary and confidential. Experimental research only — not investment advice.