



Long Game

Do enforced trading limits work — and does making them harder to lift help, or hurt?

Maksim Borisevich · ORCID 0009-0006-5309-2895 · Stage 1 Registered Report in preparation · September 2026

THE GAP

Since October 2023 the Financial Conduct Authority has required a 24-hour cooling-off period for first-time cryptoasset investors. Since 31 October 2025 the Gambling Commission has required a minimum 24-hour delay before a customer may increase a self-set limit, on the grounds that the delay allows reflection and prevents impulsive decisions.

Neither published an evidence base. The Commission's consultation response cites no pilot and no effectiveness data; the Authority's only publication on its own rule is a compliance assessment containing no behavioural measurement. A parameter written into two regulatory regimes and applied to millions of people has no published trial behind it.

THE STUDY

Fifteen to thirty retail cryptocurrency traders, ninety days each, trading their own money on their own exchange account as they already do. Five self-set constraints are enforced by a third party rather than merely displayed: a daily loss ceiling, a position-size cap, a trades-per-day cap, a post-loss cooldown and a losing-streak cooldown.

Each participant is their own control, and nobody is randomised into a no-protection arm — which would be indefensible with real money. The suspended constraint is disclosed to the participant every day it applies.

BLOCK 1 Daily loss	BLOCK 2 Position size	BLOCK 3 Trades/day	BLOCK 4 Post-loss cooldown	BLOCK 5 Streak cooldown
------------------------------	---------------------------------	------------------------------	---	--------------------------------------

18 days each · the shaded head of every block is the three-day washout, enforced but excluded from analysis · block order randomised per participant

- H1 (primary) — displacement. Suspending one constraint increases trading intensity on at least one dimension that constraint did not govern. This is risk compensation, never measured inside a single trading account across the dimensions of a tool set.
- H4 and H5 — requests and abandonment by arm. Half the participants can have a loosening applied after a 24-hour delay unless they call it off; half cannot have one at all. Both arms can ask, because the request is the observable moment at which a constraint bound and the participant wanted out.

WHAT ALREADY EXISTS

- The instrument. A Telegram bot holding a trade-only exchange credential that cannot withdraw funds and has no withdrawal function anywhere in its codebase. It refuses a key whose permissions it cannot read, never touches holdings that predate connection, and has run four months on my own account.
- The documents. Stage 1 protocol, safety case, consent form and preregistration drafted; the screening instrument diffed word by word against its published appendix.
- The audit. A 108-check acceptance protocol, run twice against the live system — once by code inspection, once by placing real orders. It found sixteen failures and twenty-two defects, nine now closed.

THE FINDING THAT CHANGED THE DESIGN

A Monte Carlo simulation under a strict behavioural null — participants whose latent appetite is drawn once and never varies by block — showed the primary hypothesis returning a false positive in 100% of replications. Because the five constraints share state, H1 as originally registered could be confirmed by arithmetic with no behaviour involved. It has been restricted to the constraint-dimension pairs that survive every parameterisation, and the calibration procedure will be registered rather than a fixed pair list, because the coupling turns out to be a property of the limits each participant chooses.

CURRENT STATUS

Everything is written; nothing has started. The system refuses to begin screening without a named responsible person to receive an escalation from the risk item, and that name is currently a placeholder. That single requirement is what is preventing recruitment.