Research Statement

Drone Bassan • McGill University, Computer Science and Economics (U2) • a16z Frontier Research Intern, Summer 2027

What I want to study

Perpetual futures are the largest market in crypto by volume, and funding is the largest recurring cash flow inside them. Every venue uses a funding payment to pull the perp price back toward spot, but each venue picks its own formula. Hyperliquid pays every hour with a one basis point interest term and a five basis point clamp on the premium. Binance pays every eight hours with a different averaging window. My question is simple to state and open in practice. What should a funding mechanism look like, and how far are the ones in use from it.

What I have already found

This fall, on McGill's student quant fund, I built the empirical side of this question. I pulled 3.5 years of hourly Hyperliquid BTC funding and the matching 8 hour Binance series through public APIs, split the data into development, validation and an untouched holdout, and asked how persistent funding is. An AR(1) fit on hourly data gives phi of 0.87, so a shock to funding is half gone in about five hours. On daily sums phi is 0.80. That means a forecast of the next week's funding collapses to the long run mean, and a trailing 7 to 30 day average beats the model as a forecaster.

The result that motivates the research is the venue gap. Over the same three years the same coin paid about 18.7 percent a year on Hyperliquid and 9.2 percent on Binance in the development window, then 7.7 percent and 3.7 percent in validation, with 0.86 weekly correlation between the two. A hedged position that stayed in the whole time earned about 14 percent a year on all capital at 3x in development and 5.5 percent in validation, after a 0.37 percent round trip cost. So funding is doing two jobs at once. It anchors the price, and it transfers wealth from longs to shorts, and the second job is twice as large on one venue as on another for no reason anyone has written down.

Why it is hard

Funding is a feedback loop. The rate depends on how traders are positioned, positioning responds to the rate, and the data an outsider can see is the equilibrium rather than the response function. The clamp, the interest term and the payment interval all shape that loop, and venues copy each other's parameters without evidence. Separating what the mechanism does from what the crowd does needs both careful empirical work across venues and a model that can be run with parameters that do not exist yet.

What I would do at Frontier

First, extend the empirical study to Bybit, dYdX and OKX so the venue comparison has more than two points, and add open interest and liquidation data to see what the crowd was doing when funding moved. Second, build a small agent based simulation calibrated to the observed persistence, where the clamp width, the interest term and the interval are inputs, and measure each one's effect on tracking error and on the carry a hedger earns. Third, write it up as a parameter recommendation that a venue could actually ship. Frontier is the right place because this is half empirical finance and half mechanism design, and because the firm's portfolio includes people who run these markets and would want the answer.

Why me

I have traded on this venue with real money and built the path orders take. Last summer I built the Android client of a trading app on Hyperliquid at Talis Trade, including the execution and market data plumbing, so I know which numbers a venue publishes and which it hides. I have the data, the persistence model and the cost model already built. I use AI tools heavily and keep a full record of how, which I think matters for research that other people need to trust.