Ten Minutes to Give It Back

A live trading journal on winning small, sizing up, and what the fills actually paid.

De’Mondre Zimmerman · November 2, 2026 · Independent working paper

At 6:20 p.m. on October 7, my model took its biggest loss since I had switched it on. I was still ahead. Ten minutes later, I had to change the tense.

My first paper on V3 started with the final win rate. Reading it back, I could see the losses but barely recognize the week that made me willing to take them. I went back to the live orders, the individual fills, and the settings I changed. One detail makes the decision to scale harder to dismiss in hindsight. The first time a loss really hurt, two bigger trades earned it back.

September 29 · Forty-six cents

I switched V3 live at 11:38 a.m. Chicago time. Its first completed position was one copper contract bought for 74 cents. It paid a dollar. The next was natural gas at 80 cents. Another dollar. By 6:30 that evening the strategy had made 46 cents. What mattered to me was that someone had actually taken the other side of both orders. I could follow my code all the way to an exchange fill and back to the money.

V3 placed resting orders on the favored side of short-dated commodity contracts. It skipped a market if the YES price had traded both at or below 45 cents and at or above 55 cents in the previous ten minutes. I wanted to avoid a favorite that had only just become the favorite. A contract bought for 90 cents earns 10 cents if it wins and loses 90 cents if it doesn’t, before fees. A long string of correct calls is quite ordinary at that price. I still needed to get paid enough when I was right.

September 30 and October 1 · The first time size helped

On Wednesday all eight positions won. Daily profit was $12.98, bringing the running total to $13.44. I was already moving the size setting around. One silver order captures the difference. I requested 25 contracts. The risk limit admitted one at 88 cents. The entire trade earned 12 cents. That day’s actual fills ranged from one contract to eleven. Looking back, I can’t draw a clean line between a fixed-size winning week and a later week when I started experimenting. The experiments had already started. So had the gains.

On Thursday I asked for 95 copper contracts. The risk controls sized the order down to 17 before submission. All 17 filled at 73 cents. All 17 lost. I lowered the setting to 65 two minutes after the loss posted. Later, 65 natural gas contracts earned $7.15 and 65 crude oil contracts earned $6.50. Together they covered the $12.41 copper loss. With an earlier gas win included, the day finished up $4.44. I had seen size rescue a day. I hadn’t yet seen it do the reverse.

01 / THE DAILY LEDGER

What I could see at the close

Starting from zero strategy profit, this is what V3 had booked each evening. A flat Saturday means that no V3 market closed in this ledger. The curve excludes deposits, other strategies, and open positions.

V3 live markets closed, Chicago dates, 2026. P&L in US dollars.

Date	Wins / losses	Day P&L	Running P&L
Sep 29	2 / 0	+$0.46	+$0.46
Sep 30	8 / 0	+$12.98	+$13.44
Oct 1	3 / 1	+$4.44	+$17.88
Oct 2	2 / 0	+$3.75	+$21.63
Oct 3	0 / 0	$0.00	+$21.63
Oct 4	2 / 0	+$5.80	+$27.43
Oct 5	6 / 0	+$24.45	+$51.88
Oct 6	7 / 0	+$20.29	+$72.17
Oct 7	2 / 2	−$82.52	−$10.35
Oct 8	1 / 1	−$21.78	−$32.13
Oct 9	1 / 1	−$6.68	−$38.82
Oct 10	0 / 0	$0.00	−$38.82
Oct 11	3 / 0	+$4.10	−$34.72
Oct 12	3 / 1	−$5.10	−$39.82

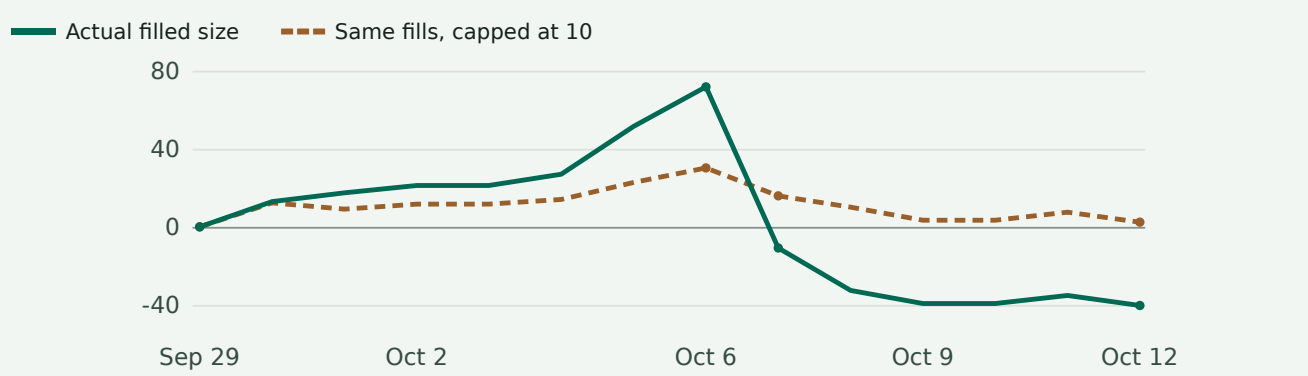


Figure 1. Cumulative profit at each calendar day’s close. The solid line omits intraday swings, including the October 7 high. The dashed line is a retrospective sizing calculation.

Values are summed before rounding. The downloadable file gives one row per market, including fractional fills. An October 9 desk review had put October 6 at −$32.71. Restoring the cash from 53 contracts closed early adds $53.00 to that settlement-only tally and gives +$20.29 for the day. The correction is already in this table.

The case for going bigger

October 2 added $3.75. October 3 had no closes. On October 4 I booked another $5.80. By Sunday night I had 17 winning markets out of 18 and $27.43 of realized profit. It was a small amount of money. It was also my money, and I could trace it back to the orders.

Monday made the argument stronger. Six wins across natural gas, copper and gold added $24.45, nearly as much as the whole run had earned through Sunday. Tuesday added $20.29. Two positions that day had been sold in whole or in part before expiry. Counting the sales, their fees, and the remaining settlement together, those two earned $8.00. Reading only their settlement records would have called them losses.

Tuesday evening’s score was 30 winning markets out of 31. The strategy had closed seven trading days in profit. My target for the pilot was an average of at least $30 a day. Monday alone had come within six dollars of it. The missing observation was how it would behave across several large losing fills. I had exactly one loss to learn that from.

October 7 · What happened between the size setting and the fill

At 3:11 p.m. the desk sent a 39-contract natural gas order. A minute later I authorized a capacity pilot with a fivefold sizing multiplier and a ceiling of 155 contracts. The gas order was already in. It earned $3.90. A later 60-contract crude oil fill earned $6.60. At 4:15 p.m., realized profit since launch stood at $82.67, its high for the study. Two other orders in copper, for 60 and 61 contracts, were canceled without a fill. The number I dialed in and the inventory I actually carried were two different things.

At 6:11 the desk bought 61 natural gas contracts for 86 cents apiece. If they had won, they would have added $8.54. The ledger booked the zero payout at 6:20:47. The loss was $52.46. I had $30.21 of cumulative profit left.

Twelve seconds after that close, the desk submitted the next order. Gold. Fifty-two contracts at 78 cents. At 6:30:17 the ledger booked another zero payout. The new running figure was **−$10.35**. Two fills had cost $93.02. Nine minutes and thirty seconds separated the two recorded closes.

I didn’t need a theory about gold and gas moving together to explain the reversal. Either loss fit inside the profit I had built. Together they were bigger than it. That afternoon I had also put in an 89-cent entry ceiling. Gas at 86 cents and gold at 78 both qualified. So did the next day’s losing gas trade at 80. The cap may have kept other orders off the book, but it let these three straight through. All three filled at the intended maker price.

The same calls with fewer contracts

Thursday’s first close was another natural gas loss, $35.20 on 44 contracts. From Wednesday’s high to Thursday’s low, three consecutive losing markets took $128.22, more than all 40 winners in the entire study would earn together ($116.51). Silver earned $13.42 later that evening. At 9:15 p.m. on October 8 I rolled back the fivefold pilot and set V3 to ten contracts. Friday still lost $6.68. Sunday’s three wins recovered $4.10. Monday’s three crude oil wins earned $3.60 between them, then one ten-contract gas loss cost $8.70. I ended the October 12 window down $39.82. Reducing size reduced the bill. It didn’t fix the arithmetic of buying favorites.

I tested the explanation I wanted to believe. What if I had kept every market I actually traded, entered at the same prices, taken the same exits, and capped each filled position at ten contracts? The comparison has to charge the smaller strategy for the profit it would have given up on the winning trades, too.

One cap, applied to winners and losers alike. Period P&L, dollars.

Period	Actual	Capped at 10
Sep 29 – Oct 6	+$72.17	+$30.65
Oct 7 – 8	−$104.30	−$20.10
Oct 9 – 12	−$7.68	−$7.68
Full window	**−$39.82**	**+$2.86**

The full-window totals use unrounded values. A ten-contract cap leaves smaller positions at their actual size. This calculation cannot predict a smaller order’s queue position or fill.

The cap gives up $41.52 of the early profit to avoid $84.20 of the two-day loss. It finishes with less than three dollars. That is a different ending to my story and a very thin basis for funding the next one.

What I can defend

Thirty of my first 31 markets made money. Only ten of the next 15 did. Larger positions paid some of the early gains and amplified the reversal. Over October 7 and 8, my three winners actually averaged slightly more contracts than my three losers, 53.3 versus 52.3. The average winner earned $7.97. The average loser cost $42.74. I can’t establish from 46 traded markets that the original rule had a durable edge, or that scaling alone caused the losses. I chose the size as I watched performance. The set of markets, entry prices, and counterparties also changed.

For the next run I would write down the rule, size cap, and end date before the first order. I would retain every eligible signal, including rejects and orders nobody filled, and reconcile each position after it closes. I’d want to see whether the next block of unseen closing times can pay for its losses at a fixed size before I ask what five times that size might earn.

How to check the account

Sample and clock

This retrospective covers September 29 through October 12, 2026, inclusive, in America/Chicago time. It uses completed, live, exchange-linked V3 positions. I excluded shadow and demo orders, the separately named V3S strategy, and the rest of the account. The narrative reconstructs events from the desk records. The first-person judgments are my interpretation in hindsight.

Unit of observation

The attribution ledger has 49 closing rows and 46 distinct markets in this window. On October 6 a single WTI position closed in three early-sale lots and one settlement lot. I combine all four so it counts as one win. A natural gas position sold fully before expiry also belongs in the sample. I assign each market to the Chicago date of its last close and use the sum of its recorded net cash outcomes, including recorded fees. Forty of the 46 markets made money. I matched 80 distinct exchange fills to the settlement records. All 46 market cash flows reconcile to the attribution ledger. Total recorded exchange fees were $0.2465. Net here means after those fees, before overhead and taxes.

Why the earlier paper says 88%

At the cutoff underlying Working Paper 02, 37 of 42 *closing rows* were positive and their net was −$38.82. That total is reproducible through the October 9 closes, before V3 next closed a market on October 11 in the evening. Grouping those same rows gives 34 winning markets out of 39 and the same dollars. This edition adds the evening of October 11 and the October 12 session. The change in win rate reflects both the counting convention and the longer window.

Sizing and execution checks

For each market with filled quantity q and recorded net P&L p, the dashed curve adds $p \times \min(q, 10) / q$. It scales all exit lots and fees proportionally. It holds the market selection, prices, and outcomes fixed, ignores fee-rounding changes and possible changes in order priority, and assumes the smaller size would have filled. Its purpose is to isolate dollar exposure on the historical fills.

The earlier offline replay gives a separate check on the signal. A 15-second delay on held-out recordings yielded 137 modeled fills across 93 closing times. Giving each close equal weight produced −1.17 cents per contract with a 3.71-cent standard error. That uncertainty is one reason I treat the profitable first week as a result worth studying and the next week as an open question.

Data and references

1. Zimmër Trading. Live order, fill, settlement, position-attribution, and operator control records. To verify a CSV row, add residual settlement payout and $1 per offsetting YES/NO pair, then subtract equivalent gross acquisition cost and recorded fees. Dated September 30, October 6 and October 9 desk reviews and the operator log document the settings. 46-market CSV.
2. De’Mondre Zimmerman (2026). How an 88% Win Rate Lost Money. Zimmër Trading Working Paper 02. See the out-of-sample execution test.
3. Constantin Bürgi, Wanying Deng and Karl Whelan (2026). Makers and Takers: The Economics of the Kalshi Prediction Market. CESifo Working Paper 12122. Independent context on favorite prices and returns to different execution roles.

The author operated the desk. AI tools assisted with the ledger reconstruction, drafting, and editing. This paper has not undergone external peer review.