

Pre-declaration — overnight hold, conditioned on volatility

Sealed before recording began. Amended once, the same day, before any observation existed — the amendment is disclosed in section 4.

A genuine seal document from my own programme, included because it shows the method at the moment that matters: *before* any result exists. Note in particular section 4, which states the known cost of the design in advance, and section 5, which records what the candidate is *not* allowed to be used to argue.

1. THE DIAGNOSTIC THAT PRODUCED THE HYPOTHESIS

An unconditional overnight hold earns roughly \$2,613 a year on one contract against a \$9,085 maximum drawdown — three and a half years of profit held as drawdown, which is awkward for a small account. The question was whether that drawdown is structural or concentrated.

Splitting its 2,637 observations by a volatility state known **before** each entry (the prior 20-session average range):

VOLATILITY AT ENTRY	N	NET	PER TRADE
Lowest fifth	528	+\$2,156	+\$4.08
Second	527	+\$1,176	+\$2.23
Middle	527	+\$10,888	+\$20.66
Fourth	527	+\$12,766	+\$24.22
Highest fifth	528	+\$259	+\$0.49

The top fifth returns 1% of the profit and carries an \$8,291 drawdown on its own. The remaining 80% earns \$12.80 per observation — twenty-six times the rate.

2. THE RULE, STATED IN FULL

The unconditional rule, but stand aside when the prior-20 average range is in the top fifth of recent history. Everything else identical: long from the 16:00 close to the next 09:30 open, no stop, no target, one micro contract, risk unit $1.00 \times$ average range.

DEVELOPMENT	N	NET	MAX DD	RETURN / DD	WIN %
Unconditional	2,637	+\$27,245	\$9,085	3.00	55.3
Conditioned	2,108	+\$26,960	\$4,120	6.54	56.4

These development figures are not evidence. They are the reason for the hypothesis, stated so that the forward record can be compared against them.

3. THRESHOLD DEFINITION, AND THE CONSTRAINT THAT SHAPED IT

The threshold is a rolling 80th percentile over the prior 500 sessions, carried by a development baseline frozen at seal time and extended by forward observations as they accrue. Every value in that baseline was known before the seal, so there is no lookahead, and the threshold adapts rather than decaying.

A shorter lookback does not substitute: return-over-drawdown falls from 7.16 at 500 sessions to 3.70 at 250, 3.38 at 120, and 2.78 at 60 — worse than applying no condition at all. A short baseline cannot establish what "extreme" means.

4. AMENDMENT, DISCLOSED

The first version of this seal used a fixed threshold — the development 80th percentile, as an absolute level. Current volatility runs well above that level, so on release it would have excluded 92% of the forward period, recording 6 observations where the unconditional rule recorded 74.

That was corrected to the rolling definition above **before any observation was recorded under it**. A candidate that knowingly records the wrong thing is worthless, and the correct response is to rebuild and disclose rather than to leave it accumulating a meaningless record. The amendment is recorded here permanently rather than in a revision that replaces the original text.

5. STATED COSTS AND LIMITATIONS, BEFORE THE FIRST OBSERVATION

- **This adds a parameter to a rule whose main virtue was having none.** That is a real cost. Both versions record side by side precisely so the forward record shows whether the condition earns its parameter.
- **The quintile split was chosen from the data.** A natural five-way cut is a weak form of searching, but it is still searching, and the improvement above is in-sample.
- **A related filter has been refuted before** in this programme — excluding a subset by risk size discarded that subset's winners along with its losers. This case differs in that the excluded bucket has almost no winners to discard (\$259 across 528 observations). That does not make it correct; it makes it a different failure mode.
- **Review point: 250 observations, or 2027-06-30,** whichever comes first. Fixed here so it cannot be moved once results exist.

6. WHAT THIS CANDIDATE MAY NOT BE USED TO ARGUE

No protected-period evidence was used in its construction — the diagnostic ran entirely on development data. It does not reopen any held question elsewhere in the book, and no existing candidate has been modified, stopped or de-listed as a result of it. The unconditional version continues recording unchanged.