

SYSTEM EXPLAINER

Gann Paper Trader

A complete guide to the course-derived research model, the engineering additions, the autonomous Alpaca paper-trading boundary, and the ticker analysis lab.

COURSE	Rules or concepts explicitly taught or demonstrated in the supplied course.
SYSTEM ADDITION	Deterministic implementation choices needed to turn the lessons into software.
ENGINEERING SAFEGUARD	Statistical, risk, paper-only, audit, and deployment controls added for safety.
COURSE BOUNDARY	Material that is missing, interpretive, undisclosed, or intentionally not automated.

Prepared 29 July 2026 | Paper trading only | Not financial advice

CHAPTER 01

Executive summary

The system is an auditable research and paper-execution pipeline. It converts selected market pivots into lunar and planetary cycle hypotheses, requires time and polarity agreement, checks market-specific price geometry, waits for completed-bar structure confirmation, and only then creates a paper-order candidate.

Most important distinction: implementation completeness is not predictive validity. A model can represent every deterministic course step and still remain statistically unqualified. The qualification gates are engineering safeguards, not promises made by the course.

The system has four products:

- Research model: fits and freezes a market-specific lunar, inversion, COR, AOR, timing, and execution protocol.
- Analysis lab: lets a user choose a ticker and past/future display periods without changing that ticker's locked settings.
- Operations dashboard: explains model status, forecasts, portfolio, decisions, risk gates, and system health.
- Paper execution worker: evaluates completed daily bars and may submit only Alpaca paper orders inside the authorized mode and risk envelope.

The repository contains no real-money routing path. Paper results still omit important live-market effects and must not be treated as expected live results.

CHAPTER 02

Evidence base and source hierarchy

The explanation is based on the supplied course and the implemented repository.

Source	Role	What it establishes
Welcome and Lessons 1-15	Course	Narrative workflow, research priorities, examples, execution philosophy, and assignments.
Moon Data .xlsx / Lunar workbase.ods	Course data	Geocentric lunar observations and the multi-history calculation workbase.
Cycles.xlsx / Astro homework.xlsx	Course data	Heliocentric longitudes, continuous values, pair/triple expressions, and research examples.
Inversions.pdf / Angle of Repetition.docx	Course handouts	Inversion concepts and price-angle work.
COURSE_MODEL_V2.md	Implementation	Frozen algorithm, validation order, execution selection, and qualification gates.
COURSE_COVERAGE_AUDIT.md	Audit	Lesson-to-code reconciliation and explicit automation boundaries.
Code and tests	Implementation evidence	Actual behavior of research, strategy, risk, broker, storage, API, and web pages.

When a lesson is interpretive or omits a numeric rule, the system must expose the boundary. It must not silently invent a value and label it a course rule.

CHAPTER 03

The course method in one view

Stage	Course question	Output
1. Market selection	Which market has usable history and a meaningful reference?	Market profile and dataset
2. Pivot map	Which swings represent the market's net force?	Alternating high/low events
3. Astronomical join	Where were the Moon and planets at each event?	Continuous angular series
4. Lunar decoding	Which angular cycle connects the pivots?	Cycle length, histories, inversion map
5. COR decoding	Which heliocentric expression repeats?	Planet expression, cycle template, polarity
6. Forecasting	Where do prior projections converge next?	Time window, expected high/low, orb
7. Price work	Which market-specific geometry aligns?	AOR and price zones
8. Execution	Does price confirm and fit the risk plan?	Trade, wait, reject, or halt decision

Course-derived core concepts

- Use geocentric Moon data for the lunar model and generally heliocentric data for planetary COR research.
- Unwrap 0-360 degree longitudes into continuous coordinates before measuring multi-revolution cycles.
- Retain multiple historical starts and project the selected cycle repeatedly into the future.
- Use the latest three or four valid histories to form the next lunar forecast.
- Track expected high/low polarity and account for inversion behavior.
- Seek confluence between time, price, and market structure rather than trading a date alone.

CHAPTER 04

Short, medium, and long horizons

Finding: the course does not require exactly three predictions in every calendar month.

Horizon in the course guide	Candidate astronomical bodies	What it means
Short	Moon or Mercury	Faster research layer; not a mandatory monthly signal.
Medium	Venus or Mars	Intermediate research layer; frequency depends on the fitted market relationship.
Long	Jupiter or Saturn	Slower research layer; naturally produces fewer events.

The course also says to repeat research on daily, weekly, and monthly charts without mixing horizons in one result. These are analysis scales, not a quota.

Why the current system may show zero or one window in a month

- The active model fits one selected lunar lineage and one selected COR expression for the ticker.
- A date survives only when lunar and COR timing and polarity agree within the frozen tolerance.
- Optional time-count confluence may remove another candidate.
- Price geometry and structure confirmation govern execution eligibility inside a surviving window.
- Sparse output is therefore expected. Adding filler dates to force three per month would contradict the confluence rule and corrupt qualification statistics.

Current implementation gap: it does not run three simultaneous, separately qualified short-, medium-, and long-horizon models. Building those would be a new research extension. Each horizon would need its own pivots, candidate family, validation, forecast density, and qualification result. The three outputs must not be combined as if they were independent evidence.

CHAPTER 05

From pivots to a lunar forecast

Course rule

- Identify significant alternating market highs and lows.
- Join those events to lunar positions and test narrowed cycle candidates: 34, 40, 42, 44, 46, 48, and 54 degrees.
- Keep the most recent three or four valid projected histories and average their convergence.
- Alternate expected high/low polarity unless an inversion rule applies.

System addition

- A causal ATR reversal detector makes the pivot definition reproducible.
- Each pivot stores its extreme date and later confirmation date, preventing look-ahead leakage.
- The ATR period, reversal size, minimum bar separation, and forecast orb are frozen per ticker.
- Chronological inner folds select a stable candidate rather than the best full-history fit.

Historical pivot markers are not historical predictions. They are causally confirmed labels used to train or score the model. A displayed forecast window is a prospective model output.

CHAPTER 06

Inversions and planetary COR

Course-derived behavior

- The lunar time cycle can continue while expected high/low polarity flips in an inversion region.
- COR is researched through one-, two-, or three-body heliocentric expressions.
- Pair and triple formula signs follow the supplied workbook conventions.
- The course seeks repeated relationships around market turns and confluence with the lunar forecast.

Implementation choices

- The system fits the Aries inversion zone plus a training-frozen variable 30-degree zone and a separate aspect trigger.
- COR candidate lengths are derived from repeated 1-12 year coordinate differences.
- Recurring pivot phases form a multi-event template; the code does not assume exactly one event per 360 degrees.
- Only lunar events with nearby COR agreement in time and polarity survive.

Undisclosed boundary

The course says COR has two fixed inversion zones but does not supply their numeric boundaries. The system does not optimize substitute zones after seeing results. It skips lunar/COR polarity disagreements in that uncertain area.

CHAPTER 07

Price geometry, time counts, and execution

Component	Origin	Automation status
Square-of-9 / AOR	Course	Market-specific angle is fitted on development data; live price must be near a frozen AOR level.
Natural squares and Fibonacci time	Course	25 and 21 starts; calendar and trading-bar variants are separate hypotheses.
Gann grid	Course boundary	Research only because it requires a manually frozen swing and chart scale.
Wheel 12/24/36	Course boundary	Deterministic calculator, but wheel selection remains a declared research trial.
Planetary price lines	Course boundary	Candidate levels only, not automatic evidence.
Break of structure	Course-derived execution	Completed daily close must confirm direction inside the active window.
Stops, targets, holding period	Course plus system	Selected on validation data from deployable candidates with explicit costs.

Price confluence is an execution condition. It must not remove a failed timing forecast from the statistical denominator; otherwise timing accuracy would be overstated.

CHAPTER 08

Ticker-specific settings and the analysis lab

Each ticker has exactly one locked settings profile. Selecting a different historical or future display period filters the output only; it does not retune the ticker. A ticker without a profile is rejected instead of borrowing another ticker's settings.

Ticker	Profile	ATR	Reversal	Separation	Orb	Selection
SPY	spy-course-v1	14	1.5 ATR	3 bars	+/-2 days	Independent SPY course fit; outer pivot detector frozen explicitly
QQQ	qqq-course-v1	14	1.5 ATR	3 bars	+/-2 days	Independent QQQ course fit; outer pivot detector frozen explicitly
IWM	iwm-course-v1	14	1.5 ATR	3 bars	+/-2 days	Independent IWM course fit; outer pivot detector frozen explicitly
DIA	dia-course-v1	14	1.5 ATR	3 bars	+/-2 days	Independent DIA course fit; outer pivot detector frozen explicitly
EUR/USD	eurusd-grid-20260729-v1	14	1.0 ATR	2 bars	+/-3 days	Selected from 36 EUR/USD-only configurations on validation data
GBP/USD	gbpusd-grid-20260728-v1	10	1.5 ATR	2 bars	+/-2 days	Selected from 36 GBP/USD-only configurations on validation data

Page behavior

- The profile list comes from a versioned configuration file.
- Yahoo daily candles are used for interactive equity and forex research; the unfinished current candle is excluded.
- The full course model is cached per ticker and daily data state. Date-range changes are fast filtered views.
- The page shows price history, causal pivots, future windows, monthly forecast density, locked settings, and out-of-sample qualification.
- The model hash includes the frozen pivot-detection settings, making a changed profile a new artifact.

CHAPTER 09

Chronological validation and qualification

Validation protocol

- Market history is split chronologically into 60% training, 20% validation, and 20% untouched final testing.
- Candidate stability is ranked across inner chronological folds.
- Only selected parameters are refit through the end of the 80% development period.
- The final 20% is scored once and is not reopened for selection.
- The experiment ledger tracks timing and execution candidates and produces an adjusted timing p-value and approximate PBO.

Mandatory autonomous qualification gates

Gate	Threshold	Origin
Untouched forecasts	At least 30	Engineering safeguard
Time hit rate	At least 55%	Engineering safeguard
Polarity accuracy	At least 52%	Engineering safeguard
Baseline lift	At least +5 percentage points	Engineering safeguard
Net expectancy	Positive after assumed costs	Engineering safeguard
Probability of backtest overfit	No more than 50%	Engineering safeguard
Adjusted timing p-value	No more than 0.05	Engineering safeguard
Leakage, hash, immutable protocol	All required	Engineering safeguard

The course defines the modeling workflow but does not provide these numerical qualification cutoffs. The thresholds were added to stop autonomous execution from treating an attractive in-sample fit as validated evidence.

CHAPTER 10

How to read the key metrics

Metric	Meaning	Common misreading
Time hit	Share of forecast windows close enough to a causally confirmed pivot.	It is not directional accuracy and does not mean profitable trades.
Polarity accuracy	Correct high/low label among scored matches.	A high score on few matches can be unstable.
Baseline lift	Improvement over a frozen naive timing baseline.	A positive raw hit rate can still have negative lift.
Forecast density	Number of windows per month or year.	More forecasts are not automatically better.
Net expectancy	Average R after the separately defined execution rule and costs.	A positive value on three trades is weak evidence.
Adjusted p-value	Timing evidence after multiple testing adjustment.	It is not the probability that the model is true.
PBO	Approximate probability that selection is overfit.	It does not remove all data-mining risk.

Why a model remains unqualified after code improvements

Engineering improvements can fix leakage, make rules complete, increase transparency, and implement missing course calculations. They cannot force untouched market data to validate the hypothesis. Percentages change only when the ticker's frozen settings, market data, or evaluation sample changes. A model should remain unqualified when the evidence fails.

CHAPTER 11

Decision and paper-order lifecycle

Stage	Condition	Outcome
WAIT	No active forecast window.	No entry evaluation.
WATCH	Forecast is active but confirmation is absent.	Continue observing completed bars.
REJECT	Qualification, polarity, price, structure, risk, or policy gate fails.	No order.
SIGNAL	Model and execution conditions pass.	Create a sized paper-order intent.
ORDER	Idempotency and reconciliation pass.	Submit Alpaca paper bracket order.
POSITION	Paper broker reports exposure.	Monitor holding, exits, risk, and reconciliation.
HALT	Kill switch, daily loss, broker block, or integrity condition triggers.	Block entries and apply halt policy.

Order safety additions

- Alpaca client construction is hard-coded to paper mode.
- Deterministic client order IDs prevent duplicate submission.
- Startup and periodic reconciliation compare local intent with broker state.
- The worker uses completed daily bars, not an unfinished candle.
- The kill switch and database halt state require explicit recovery acknowledgment.

CHAPTER 12

Risk policy and autonomous boundaries

The bounded multi-ticker pilot policy currently uses:

Control	Pilot value	Reason
Risk per position	0.10% of equity	Limit damage while models are unqualified.
Concurrent positions	4	Bound portfolio breadth.
Symbol notional cap	5%	Prevent single-name concentration.
Gross exposure cap	20%	Cap total simulated exposure.
Daily loss halt	0.50%	Stop autonomous entries after a bad day.
Reward/risk target	3:1	Frozen pilot execution parameter.
Short entries	Allowed in paper pilot	Test bearish HIGH forecasts.
Pyramiding	Blocked	Outside authorized risk scope.
Extended hours	Blocked	Avoid thin and different execution conditions.

An explicit bounded pilot may test an unqualified candidate in paper trading, but this exception does not make the candidate qualified and must remain visibly labeled.

CHAPTER 13

Dashboard, health, audit, and deployment

Operations dashboard

- Shows each ticker independently: stage, latest decision, active/next forecast, qualification, position, and open orders.
- Shows portfolio equity, buying power, positions, risk envelope, worker heartbeat, reconciliation, and halt state.
- Explains unusual terms through inline notes and a glossary.
- Separates active execution artifacts from research candidates.

Audit and persistence

- SQLite stores decisions, order intents, audit events, state, heartbeats, and reconciliation results.
- Frozen model protocols contain input hashes, split boundaries, candidates, metrics, and model hash.
- Content-addressed experiment files prevent silent overwrite of prior trials.

Deployment additions

- Docker isolates the application and state volume.
- The website binds locally behind a reverse proxy and is published under amr.alafaleq.cloud.
- Repository polling, health checking, rollback, subdomain routing, and scheduled maintenance are deployment operations, not course rules.
- Credentials must remain in server environment configuration and are redacted from APIs and the dashboard.

CHAPTER 14

What is not automated or not proven

Area	Status	Why
COR numeric inversion zones	Conservative boundary	The supplied course does not disclose exact boundaries.
Manual natal/reference event	Research judgment	Market birth definitions are interpretive.
Gann grid scaling	Research only	Requires a frozen chart scale and swing.
Wheel selection	Research only	Several deterministic ladders exist; choosing one needs a declared trial.
Retracement / break-even / trailing	Research only	Exact live child-order replacement semantics are not implemented.
Pyramiding	Excluded	Expands exposure beyond approved policy.
Three simultaneous horizon models	Not implemented	Would require three separately specified and validated model families.
Causal astronomical claim	Not established	The course examples are not an independently audited prospective track record.
Live profitability	Not established	Paper fills do not reproduce full market impact, liquidity, latency, and fees.

The correct reading is: this is a transparent experimental system. It can generate hypotheses, measure them honestly, and paper-test bounded execution. It cannot guarantee a turn, price, or profit.

CHAPTER 15

Operational checklist

Before trusting an analysis page result

- Confirm the selected ticker and locked profile ID/hash.
- Confirm the last completed candle date and requested past/future ranges.
- Distinguish historical pivots from prospective forecast windows.
- Read monthly forecast density; zero is a legitimate result.
- Check qualification and sample size before interpreting confidence.

Before allowing a paper order

- Confirm paper-only mode, paper credentials, market clock, worker heartbeat, and reconciliation.
- Confirm the correct per-ticker frozen model artifact is active.
- Confirm a forecast is active and price, structure, and risk gates pass.
- Confirm exposure, daily loss, order duplication, and kill-switch gates pass.

After a forecast window closes

- Score timing and polarity using the predeclared rule.
- Record missed pivots and false-positive windows, not only successful examples.
- Do not retune using the untouched final period and call the new result out-of-sample.
- A changed rule or ticker profile is a new model and must receive a new hash and experiment.

CHAPTER 16

Glossary and final synthesis

Term	Plain-language meaning
AOR	Angle of Repetition: a market-specific price/planetary angle observed at repeated turns.
COR	Cycle of Repetition: a selected heliocentric planet expression associated with market turns.
Continuous longitude	Angular position with completed 360-degree revolutions accumulated instead of reset.
Confluence	Independent model conditions pointing to the same time, polarity, or price area.
Forecast window	A bounded watch period around a projected central date; not an order.
Inversion	A region where expected high/low polarity flips while timing continues.
Orb	Allowed tolerance around an expected time or angular event.
Pivot confirmation	The later date on which a reversal makes an earlier extreme knowable.
R	Profit or loss expressed in units of the initial trade risk.
Time hit	A forecast window that lands sufficiently close to a qualifying confirmed pivot.

Final synthesis

The course contributes the market-cycle hypothesis, the astronomical coordinate work, lunar and COR histories, inversion concepts, time/price geometry, and confirmation philosophy. The software adds causal definitions, frozen per-ticker profiles, chronological selection, statistical qualification, realistic paper-cost modeling, risk controls, broker reconciliation, audit trails, web explanations, caching, deployment, and paper-only enforcement.

The system should produce as many or as few forecasts as the frozen confluence rules justify. It should never manufacture three dates per month merely to fill short-, medium-, and long-term labels. If separate horizon models are built later, each must stand on its own evidence.

End of guide