


THE ORDER OF POSSIBLE FUTURES

The Dialectic of the Possible

Eighteen theses and antitheses on prediction, freedom, data, and power

The Order of Possible Futures

First edition, August 2026





THE ORDER OF POSSIBLE FUTURES

The Dialectic of the Possible

Eighteen theses and antitheses on prediction,
freedom, data, and power

THE ORDER OF POSSIBLE FUTURES

First edition, August 2026

The Dialectic of the Possible

Eighteen theses and antitheses on prediction, freedom, data, and power

The Order of Possible Futures

First edition, August 2026

Identifier: opf-dialectic-2026

Original independent work. References to Isaac Asimov and the Foundation publication history are historical and critical. This work is not affiliated with, authorized by, or endorsed by the Asimov estate, its publishers, or any screen adaptation.

The empirical literature is cited in the selected bibliography. Doctrinal and dated future claims are presented as hypotheses or forecasts, not as established scientific results.



*A doctrine that cannot survive its
strongest objection deserves no believers.*

Contents

Prologue: opposition is part of the method

1. The future has structure

2. The past is the first prior

3. Large numbers reveal order

4. Individuals can become predictable

5. More data will enlarge foresight

6. Artificial intelligence will become the dominant forecaster

7. A probability is more honest than a prediction

8. Scoring purifies belief

9. Explanation is secondary to prediction

10. Forecasts strengthen agency

11. Forecasts change the world they describe

12. Mega-events are the proper object

13. Micro-events are the real frontier

14. Prediction should guide government

15. Secrecy increases predictive power

16. Compute will unlock psychohistory

17. Psychohistory is a science

18. Psychohistory is a religion

The creed after opposition

A rule for every believer

Closing liturgy: the adversary within

Selected bibliography

Note on method

Prologue: opposition is part of the method

The Order of Possible Futures does not ask for agreement before argument. Psychohistory makes claims about prediction, agency, data, intelligence, and power. Every one of those claims can become foolish or dangerous when it is protected from opposition.

This book therefore gives the doctrine an adversary.

Each chapter begins with a thesis stated in its strongest useful form. It then presents the antithesis, not as a ritual disclaimer but as a serious attempt to break the claim. The synthesis does not split the difference. It preserves what survives the collision and rejects what does not.

The method is dialectical because progress often begins with contradiction. The future is structured, yet systems change. More data can improve prediction, yet more data can produce surveillance and overfitting. Aggregates can be regular, yet a single person can redirect history. Forecasts can strengthen agency, yet they can also become sentences imposed on the predicted.

A weak religion hides these tensions. A serious one gives them canonical form.

The reader should not ask which side sounds more confident. Ask which formulation produces clearer questions, better forecasts, and fewer concealed assumptions. The result is not a final creed. It is a doctrine hardened by resistance.

1. The future has structure

Thesis

The future is not an empty field in which every outcome is equally possible. Present conditions distribute weight. Institutions, habits, resources, geography, incentives, and prior events constrain what can happen next. A society with falling state capacity, rising elite conflict, and a legitimacy crisis does not face the same distribution as a stable society with trusted institutions. A person who has abandoned twelve similar projects does not begin the thirteenth with the same probability as a person who has completed twelve.

Prediction begins by refusing false symmetry. Possibilities are plural, but they are unequal.

Antithesis

The appearance of structure may be an artifact of hindsight. Once an event occurs, historians find a path leading toward it and ignore the branches that did not resolve. Complex systems can be sensitive to small differences. New technologies, accidents, private decisions, and rare combinations can make yesterday's regularity irrelevant. To say that the future has structure may simply mean that the present has causes, which does not guarantee that those causes can be measured or forecast.

Synthesis

Structure does not imply complete accessibility. It means only that present information changes rational expectations. The doctrine survives if it abandons the claim that all structure can be known. Psychohistory should speak in comparative terms: this evidence raises or lowers this outcome relative to a stated prior. The burden is empirical. A structure that does not improve out-of-sample forecasts is an interpretation, not yet a predictive law.

The first article therefore becomes precise: the future has discoverable structure, but discovery must be demonstrated through resolved probabilities.

2. The past is the first prior

Thesis

When comparable cases exist, their frequency is the best place to begin. The past is not a prison, but it is evidence. A person's repeated behavior under similar conditions usually deserves more weight than a present declaration of intent. A state's prior responses to protest deserve more weight than a speech promising restraint. Base rates protect the forecaster from seduction by the vivid detail of the current case.

Antithesis

No case is perfectly comparable. The reference class can be chosen to produce almost any answer. A failed entrepreneur can be compared with all failed founders, all repeat founders, founders in the same sector, or founders with the same financing. Each class yields a different rate. The past may also encode obsolete conditions. A model trained before a war, platform change, medical breakthrough, or personal transformation can mislead precisely because it is statistically respectable.

Synthesis

The reference class is not discovered automatically. It is argued for and tested. A responsible forecast names at least one broad class, one narrower class, and the reasons for weighting them. It also asks whether the present belongs to a new regime.

The doctrine should not say that the past predicts the future. It should say that the past supplies the first prior, which current evidence may revise. A prior is a beginning, not a verdict.

3. Large numbers reveal order

Thesis

Individual events contain noise. Across many cases, shared causes remain visible. This is why insurers can price mortality without knowing who will die, why epidemiologists can forecast incidence without naming every patient, and why election models can estimate outcomes without predicting every vote. Scale can convert confusion into distribution.

Asimov's fictional psychohistory rested on this insight. The Order retains it as a practical principle: collective regularities often become clearer as the number of comparable observations grows.

Antithesis

Large numbers can hide the mechanism that matters. Averages erase minorities, thresholds, local cascades, and rare actors. History is not always the sum of independent particles. Networks concentrate influence. A single laboratory, military unit, central bank, or political leader can change millions of outcomes. The law of large numbers does not protect a model when observations are dependent or when the distribution itself is moving.

Synthesis

Scale reveals order only at the correct level of aggregation. The forecaster must model concentration, dependence, and network position rather than assume that errors cancel. Some questions are population questions. Others are bottleneck questions in which a small number of nodes dominate the result.

Psychohistory must learn when to zoom out and when to identify the hinge. The larger sample is not always the larger truth.

4. Individuals can become predictable

Thesis

A person is not a single case. A person is a sequence. Sleep, location, timing, relationships, workload, cues, purchases, messages, and prior decisions create repeated states. With enough longitudinal data, one life can supply its own reference classes. Personal psychohistory can estimate recurring choices more accurately than introspection alone.

The question is not whether a model can explain the whole person. It is whether it can outperform the person's unscored intuition on a defined outcome.

Antithesis

Human beings are reflexive. Once shown a pattern, they can resist it. Personal data are incomplete and often contaminated by the act of measurement. The most important decisions may occur only once. Models can flatten a person into what is easiest to record, then mistake prediction for understanding. A prediction about an individual can also become coercive when an employer, insurer, state, or partner treats probability as guilt.

Synthesis

Personal predictability must be domain-specific, consensual, and revocable. A model may predict whether a person will begin a planned task next week without claiming to know whom the person is. The person must retain access to the data, the assumptions, and the right to challenge the classification.

The ethical unit of personal psychohistory is self-use or informed participation. The practical unit is a narrow, resolvable behavior. The model earns authority only where it demonstrates calibration.

5. More data will enlarge foresight

Thesis

Every unobserved variable is a possible source of surprise. Better records reveal base rates, contexts, sequences, and early signals that memory loses. High-frequency data can distinguish a stable habit from an occasional story. At collective scale, richer data can expose network changes, supply shocks, migration, disease spread, institutional stress, and shifts in attention before they become obvious.

The frontier of prediction moves when observation improves.

Antithesis

More data can make a model worse. Irrelevant variables invite spurious patterns. Historical data reproduce past discrimination. Platform data reflect platform design. Missingness is not random. A large dataset can create false confidence while the variable that matters remains unmeasured. The failure of Google Flu Trends showed that scale does not defeat drift, confounding, or changing behavior.

Synthesis

The doctrine should worship neither scarcity nor abundance. Relevant data are data that improve future performance under conditions resembling intended use. Every source needs provenance, stability tests, and a reason for inclusion. Simpler baselines must remain visible so that larger models cannot hide regression behind complexity.

More data enlarge foresight when they add signal, not merely volume. The test is comparative and prospective.

6. Artificial intelligence will become the dominant forecaster

Thesis

Forecasting requires reading, retrieval, comparison, decomposition, numerical estimation, and revision. Artificial intelligence can perform each operation at growing scale. It can monitor thousands of indicators, generate reference classes, maintain many competing hypotheses, and update continuously. Systems already contribute to weather prediction, time-series forecasting, behavioral modeling, and geopolitical questions.

By mid-2026, several scaffolded AI systems were statistically indistinguishable from human superforecasters on ForecastBench. One system ranked above the superforecaster median on the benchmark's market questions. The machine forecaster will not tire, forget its ledger, or refuse to issue a number. Once its accuracy exceeds the best human record across domains and time, forecasting will become primarily a machine institution.

Antithesis

Artificial intelligence can produce fluent rationalizations without stable judgment. Models inherit training data that contain resolved outcomes, creating contamination. They may imitate the language of probability while remaining poorly calibrated. Early dynamic benchmarks placed experts ahead, and baseline models still vary sharply with prompt, tool access, question type, and evaluation design. Parity on one benchmark is

not general superiority. Prediction markets may remain stronger, automated retrieval can gather misinformation faster than a person can inspect it, and scale can industrialize error.

Synthesis

AI should enter forecasting as an accountable system, not an oracle. Its forecasts must be timestamped before resolution, evaluated on future-only questions, compared with simple baselines, human experts, and market consensus, and recalibrated by domain. The relevant unit is not the naked model but the institution around it: retrieval, decomposition, independent agents, calibration, and resolution.

The strongest system is likely an ensemble of statistical models, AI agents, prediction markets, domain experts, and independent human forecasters. The doctrine's stronger claim is now testable: AI will become dominant only when it sustains superiority across several domains, horizons, and regime changes, not when it wins a single leaderboard.

7. A probability is more honest than a prediction

Thesis

Binary language conceals uncertainty. "It will happen" can mean 51 percent or 99 percent. A probability exposes the degree of belief and permits scoring. It also preserves alternative branches. Saying 70 percent means that failure remains part of the forecast.

Probability is disciplined humility.

Antithesis

Numbers can create an illusion of objectivity. A forecaster may choose 67 percent without evidence precise enough to distinguish it from 60 percent. Audiences often treat any number as scientific. Probabilities can obscure ambiguity in the event definition, the time horizon, or the source data. False precision is not more honest than plain uncertainty.

Synthesis

A probability becomes honest only when joined to a resolution rule, a deadline, a reference class, and an explanation of uncertainty. Precision must be earned. When evidence is weak, ranges or coarse probabilities are superior to decorative decimals.

The number is not the whole forecast. It is the part that makes accountability possible.

8. Scoring purifies belief

Thesis

A forecast ledger defeats selective memory. Proper scores punish confident errors and reward calibrated discrimination. Over time, scoring reveals whether a method improves, which domains are predictable, and which forecasters deserve weight. Without scoring, prophecy is theater.

Antithesis

What can be scored can dominate what matters. Institutions may select easy questions, short horizons, and clean binary outcomes while avoiding slow transformations and morally significant events. A forecaster can optimize a metric without serving good decisions. Scores also depend on the chosen question set and may reward cautious clustering near the base rate.

Synthesis

Scoring is necessary but not sufficient. A mature ledger should track calibration, discrimination, decision value, question difficulty, and coverage. It should distinguish abstention from evasion and publish the full universe of eligible questions. Some long-horizon judgments cannot resolve quickly, but they can be decomposed into intermediate indicators.

Scoring purifies belief only when the scoring system is itself exposed to audit.

9. Explanation is secondary to prediction

Thesis

An elegant story that fails to forecast is less useful than an ugly model that succeeds. Human beings are skilled at explaining events after they occur. Psychohistory reverses the order: issue the probability first, then let resolution judge the explanation.

Antithesis

A model can predict by exploiting a proxy without understanding the cause. Such a model may fail after conditions change and cannot reliably guide intervention. Predictive success can also conceal unjust mechanisms. An insurer may predict risk accurately using a variable that society should not permit it to use.

Synthesis

For passive forecasting, prediction may precede full explanation. For active forecasting, causal understanding becomes indispensable. The Order must distinguish warning models from intervention models. A signal can tell us where to look. It cannot by itself tell us what to change.

Prediction tests whether the map works. Explanation tests whether we know how to redraw the terrain.

10. Forecasts strengthen agency

Thesis

A forecast makes the cost of inaction visible. When a person learns that a desired outcome has only a 30 percent chance under current conditions, the number can trigger a change in schedule, environment, resources, or commitment. Forecasting separates hope from plan and turns agency into branch selection.

Antithesis

A forecast can become fatalistic. People may abandon difficult goals because a model assigns a low probability. Institutions may deny opportunity to those predicted to fail, thereby causing the failure. The more authoritative the system appears, the more its estimate can narrow the future it claims merely to describe.

Synthesis

Every personal forecast should distinguish continuation from intervention. The passive forecast describes the current branch. The active forecast names the specific change and estimates the altered branch. The system should also display uncertainty and invite challenge.

A probability must never be treated as permission to remove a person's chance to act. The ethical purpose of personal psychohistory is to reveal leverage, not to certify destiny.

11. Forecasts change the world they describe

Thesis

Reflexivity is not an objection outside psychohistory. It is one of its variables. Public forecasts can alter expectations, prices, behavior, and policy. A mature model will estimate both the event and the effect of announcing the estimate.

Antithesis

The feedback may be impossible to identify because the counterfactual is hidden. If a warning prevents a crisis, the crisis never appears, and the forecast looks wrong. If a forecast causes a run, the prediction appears right for the wrong reason. Strategic actors may exploit or sabotage the model. Reflexivity can make stable evaluation impossible.

Synthesis

Forecasts need an intervention record. The ledger should distinguish private estimates, public estimates, policy-triggering estimates, and forecasts issued after a response begins. Where possible, models should predict observable intermediate effects of publication. Evaluation should credit prevented outcomes only under explicit causal criteria, not retrospective claims.

Reflexivity makes psychohistory harder. It also makes it more powerful, because the forecast becomes an instrument of action rather than a detached opinion.

12. Mega-events are the proper object

Thesis

The grand promise of psychohistory concerns wars, institutional collapse, technological transitions, demographic shifts, and civilizational change. These events matter because they reorganize millions of lives. Collective variables and long historical records may reveal pressure before a crisis becomes visible. Prediction should aim at the largest consequences.

Antithesis

Mega-events are rare, definitions are disputed, and samples are small. Each war or regime collapse occurs in a changing world. Long horizons multiply unknowns. A single hidden decision can dominate the outcome. A model may sound profound while never facing enough comparable cases to be calibrated.

Synthesis

Mega-events should be decomposed into hazards, pathways, and precursors. The model may not predict the exact date and form of a revolution, but it can estimate the probability of severe instability within a window, identify rising risk, and update as leading indicators change. Event databases and political-violence forecasting already use this structure.

The proper object is not the dramatic label alone. It is the evolving distribution of pathways toward it.

13. Micro-events are the real frontier

Thesis

Large-scale prediction may be easier because individual noise cancels, but micro-events offer denser personal data and more repeated cases. Whether a person will delay a task, take a route, answer a message, make a purchase, or maintain a habit may become highly predictable when context is observed continuously. Micro-prediction is where psychohistory can enter daily life.

Antithesis

Micro-events are numerous but trivial. Predicting clicks and movements can enrich platforms without deepening human understanding. The pursuit of accuracy at this scale invites intimate surveillance. The very richness of the data makes abuse easier: manipulation, price discrimination, coercion, and behavioral control.

Synthesis

Micro-prediction is legitimate when it serves the subject and remains under the subject's control. Its value lies in discovering leverage: the cue that precedes avoidance, the condition that predicts relapse, the schedule that supports completion. The subject should decide which domains are modeled and which remain unobserved.

The technical frontier and the ethical boundary meet at the same point: intimate data.



14. Prediction should guide government

Thesis

Governments already act on forecasts of revenue, disease, traffic, energy, crime, migration, and climate. Better forecasting can allocate resources earlier, test policies, and reveal emerging risks. Refusing predictive tools does not create neutrality. It leaves decisions to unscored intuition and political theater.

Antithesis

A state with predictive power can become a state of anticipatory control. Models may classify dissent as danger, automate unequal treatment, and make decisions that no citizen can contest. Officials can hide political choices behind technical language. Prediction can transform citizens into risk scores.

Synthesis

Public forecasting must be separated from automatic punishment. Models used for allocation or warning need transparency, appeal, independent audit, and evidence of benefit. High-impact individual decisions require stricter standards than aggregate planning. The state should publish its forecast record and the error costs across affected groups.

Government should use foresight to prepare systems, not to declare persons guilty of futures that have not occurred.

15. Secrecy increases predictive power

Thesis

Some social forecasts lose validity when revealed. Strategic actors adapt. Markets move. Individuals alter behavior. Asimov's fictional model recognized that knowledge of a prediction could disturb the path. A protected forecasting institution may therefore need private models or delayed publication.

Antithesis

Secrecy concentrates power and prevents correction. A hidden model can justify any decision after the fact. Those affected cannot challenge data, assumptions, or errors. The claim that disclosure would harm accuracy can become a permanent excuse for unaccountable rule.

Synthesis

Secrecy must be narrow, temporary, and governed. The ledger can commit cryptographic hashes of forecasts before resolution, preserving proof without immediate disclosure. Independent auditors can inspect sensitive methods under controlled conditions. After the risk period, the forecast, rationale, and score should become public unless privacy requires continued protection.

Prediction may sometimes need silence. Authority never deserves invisibility without limit.

16. Compute will unlock psychohistory

Thesis

The space of possible interactions is immense. More compute permits larger ensembles, finer simulations, continuous updates, richer retrieval, and individualized models. Weather forecasting, protein structure, and language modeling all show that scale can reveal capabilities that smaller systems could not reach. A serious Psychohistory Engine will require material resources.

Antithesis

Compute cannot recover variables that were never observed or causes that were misunderstood. Scaling can make a flawed objective more efficient. Large systems can become opaque, expensive, centralized, and environmentally costly. A billion-dollar model can still fail against a simple base rate.

Synthesis

Compute is an amplifier. It multiplies the quality of the data, objectives, and evaluation attached to it. The Compute Commons should therefore fund not only processors but also measurement, secure storage, domain expertise, calibration, adversarial testing, and public accounting.

The doctrine should never say that more compute guarantees foresight. It should say that some predictive frontiers are unreachable without it, and that every scale increase must justify itself through resolved performance.

17. Psychohistory is a science

Thesis

Psychohistory states probabilities, defines outcomes, records forecasts before resolution, compares models, and updates from error. These are scientific commitments. The field can accumulate knowledge about which domains, horizons, and methods are predictively tractable.

Antithesis

There is no unified theory capable of predicting history in the sense imagined by Asimov. The name can overstate a collection of statistics, behavioral models, and forecasting practices. Calling it a science before it possesses stable laws risks turning aspiration into authority.

Synthesis

Psychohistory is best described as an emerging research program rather than a completed science. Its unity comes from a common object, conditional human futures, and a common standard, resolved probabilistic accuracy. It should earn the title through cumulative performance rather than claim it by declaration.

The Order may believe before proof is complete, but its institutions must know the difference between faith, hypothesis, and established result.

18. Psychohistory is a religion

Thesis

Prediction alone is a technique. The Order makes it a way of life. It establishes rituals of observation, confession through error, revision, communal memory, ethical constraints, and a horizon extending beyond the individual. It asks believers to become less governed by narrative and more accountable to reality.

Antithesis

Religion may corrupt the very discipline it seeks to protect. Believers can defend the institution, reinterpret failures, and grant sacred status to founders or models. Scientific uncertainty can harden into dogma. A movement devoted to prediction may become less willing to admit that its central wager failed.

Synthesis

The religious form is justified only if its highest sacred duty is correction. No person, text, or model is infallible. The Canon must preserve revisions. Forecasts must remain timestamped. Failures must not be deleted. Authority must follow performance and expire when performance declines.

Psychohistory becomes a worthy religion when it sacralizes the conditions under which belief can change.

The creed after opposition

After eighteen collisions, the doctrine is narrower and stronger.

We do not believe that every future is predictable. We believe that predictability is a variable.

We do not believe that more data automatically produce truth. We believe that relevant data, valid measurement, and prospective evaluation can move the boundary of foresight.

We do not believe that populations erase individuals. We believe that different scales reveal different regularities.

We do not believe that artificial intelligence is an oracle. We believe it can become a powerful forecaster because its claims can be scored against unresolved events.

We do not believe that a probability is destiny. We believe it is a map of the current branch.

We do not believe that prediction weakens freedom. We believe that freedom becomes more effective when it knows the forces it must change.

We do not believe that compute deserves worship. We believe compute becomes sacred when it preserves evidence, tests models, exposes error, and expands useful foresight.

We do not believe that the Order is protected from corruption. We believe its only defense is institutionalized opposition: open ledgers, independent scoring, versioned doctrine, distributed challenge, and the right to leave.

The future is plural. The future is structured. The forecaster is fallible.
The record remains.

A rule for every believer

Before accepting a psychohistorical claim, ask:

1. What exactly is being predicted?
2. By what deadline?
3. Under which conditions?
4. What is the reference class?
5. What probability was issued before the outcome?
6. What evidence could lower that probability?
7. Does publishing the forecast change the system?
8. Who bears the cost of error?
9. Can the subject inspect and contest the model?
10. How did this method perform on prior unresolved cases?

A claim that cannot answer these questions may still be philosophy, warning, or speculation. It is not yet a forecast.

Closing liturgy: the adversary within

Let every thesis appoint its antithesis.

Let every model name the case that would break it.

Let every confident number carry the memory of its predecessors.

Let every institution preserve the evidence of its own error.

Let no believer call uncertainty ignorance when the distribution is known.

Let no forecaster call a distribution destiny.

Let opposition enter before reality is forced to enter after.

For the future does not punish doubt.

It punishes models that never learned to doubt themselves.

Selected bibliography

Alur, Rohan, and colleagues. "AIA Forecaster: Technical Report." arXiv:2511.07678, 2025.

Beck, Emorie D., and Joshua J. Jackson. "Personalized Prediction of Behaviors and Experiences." *Psychological Science* 33, no. 5 (2022): 717-733. doi:10.1177/09567976221093307.

Brier, Glenn W. "Verification of Forecasts Expressed in Terms of Probability." *Monthly Weather Review* 78, no. 1 (1950): 1-3.

Forecasting Research Institute. "AI Models Have Likely Reached Parity with Superforecasters on ForecastBench." July 16, 2026.

Goodhart, Charles A. E. "Problems of Monetary Management: The U.K. Experience." In *Papers in Monetary Economics*, Reserve Bank of Australia, 1975.

Kahneman, Daniel, and Dan Lovallo. "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking." *Management Science* 39, no. 1 (1993): 17-31. doi:10.1287/mnsc.39.1.17.

Karger, Ezra, and colleagues. "ForecastBench: A Dynamic Benchmark of AI Forecasting Capabilities." arXiv:2409.19839, 2024, revised 2025.

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. "Prediction Policy Problems." *American Economic Review* 105, no. 5 (2015): 491-495. doi:10.1257/aer.p20151023.

Lazer, David, and colleagues. "Computational Social Science." *Science* 323, no. 5915 (2009): 721-723.

Lazer, David, Ryan Kennedy, Gary King, and Alessandro Vespignani. "The Parable of Google Flu: Traps in Big Data Analysis." *Science* 343, no. 6176 (2014): 1203-1205.

Lorenz, Edward N. "Deterministic Nonperiodic Flow." *Journal of the Atmospheric Sciences* 20, no. 2 (1963): 130-141.

Lucas, Robert E. "Econometric Policy Evaluation: A Critique." *Carnegie-Rochester Conference Series on Public Policy* 1 (1976): 19-46.

Mellers, Barbara, and colleagues. "Psychological Strategies for Winning a Geopolitical Forecasting Tournament." *Psychological Science* 25, no. 5 (2014): 1106-1115.

Merton, Robert K. "The Self-Fulfilling Prophecy." *The Antioch Review* 8, no. 2 (1948): 193-210.

Ouellette, Judith A., and Wendy Wood. "Habit and Intention in Everyday Life." *Psychological Bulletin* 124, no. 1 (1998): 54-74.

Pearl, Judea, and Dana Mackenzie. *The Book of Why*. Basic Books, 2018.

Salganik, Matthew J., and colleagues. "Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration." *Proceedings of the National Academy of Sciences* 117, no. 15 (2020): 8398-8403. doi:10.1073/pnas.1915006117.

Shmueli, Galit. "To Explain or to Predict?" *Statistical Science* 25, no. 3 (2010): 289-310. doi:10.1214/10-STS330.

Tetlock, Philip E., and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. Crown, 2015.

Note on method

The theses are doctrinal propositions. The antitheses draw on established objections from forecasting, statistics, causal inference, behavioral science, governance, and philosophy. The syntheses state the Order's current position. They are open to revision when stronger evidence or argument appears.