

"Equilibrium"

The paradoxical signature of competitive markets is a positive stock of open alpha opportunities

Sam Arie

24 September, 2026

(Working Paper 3, Version 1.3)

Abstract

A long-standing intuition in finance is that competition should eliminate all opportunities for alpha. This paper models a market equilibrium in which the opposite holds. Building on two companion papers showing that innovative interpretation of the public record earns alpha but, by improving prices, simultaneously erodes that opportunity, the model makes the production of new interpretations endogenous. Investors search because better interpretation has private value, and successful ideas are then disseminated and competed away. Yet the opportunity set as a whole does not dissolve because available interpretive methods remain improvable, through innovation, while new forecasting problems continually arise. When the expected reward from discovery of new methods exceeds the marginal cost of search, the resulting stationary equilibrium contains a positive stock of open opportunities for alpha even though each individual opportunity has a finite half-life and eventually becomes uneconomic to exploit. Competition therefore eliminates particular opportunities without exhausting the general set and converges to an equilibrium paradoxically characterised by a positive stock of open opportunities for alpha.

Keywords: interpretive competition; alpha; equilibrium; investment research; price discovery; forecasting.

JEL classification: G14, G11, G12, D83, D84, O31.

Statement on use of AI tools on final page.

1. From competition to equilibrium

Competition is usually expected to eliminate investment opportunities. If a profitable strategy becomes known, capital should move towards it, prices should adjust and the abnormal return should disappear. The companion paper (Arie, 2026b) develops that mechanism for interpretation of public evidence: a superior reading can earn alpha because it identifies a direction in which price can improve, while deployment and dissemination progressively reduce the remaining opportunity. Individual interpretive opportunities therefore have finite half-lives. Individual alpha opportunities are temporary.

This paper asks whether competition also exhausts the opportunity set as a whole. Investors search for better interpretations when the expected reward from discovery covers the marginal cost of beginning that work. Successful discoveries create opportunities which competition then erodes. If new forecasting problems continue to present scope for worthwhile improvement, creation and erosion can co-exist in a stationary equilibrium.¹

The model has two outcomes. If the expected reward exceeds the cost of starting research, investors search and new opportunities keep appearing. If not, nobody searches and none appear. In the active regime, let v^* be the rate at which investors discover a new interpretation and take a position on it before anyone else gets there first. Each such discovery is then an open opportunity. Write $T = E[\tau]$ for their mean economic life. The expected number of open opportunities is

$$N^* = v^* T > 0 \tag{1}$$

¹ Here we say that a problem has ‘scope’ if some interpretation of its public record would forecast the payoff better than the price does, by enough that the resulting alpha exceeds the cost of exploiting it ($\alpha_0 > c_e$).

Let $\alpha(u)$ be the remaining normalized gross alpha of an open opportunity at age u . Write $Q = E[\int_0^\tau \alpha(u) du]$ for its expected lifetime total, averaging over initial sizes and subsequent paths. The expected normalized opportunity stock is

$$O^* = v^* Q > 0. \quad (2)$$

Search decisions determine v^* . With a positive marginal cost of starting research, search is active only when the expected discovery prize V exceeds that cost, c_r . If $V \leq c_r$, the stationary open stock is zero. When search is active, finite opportunity lives are compatible with a positive expected stock, because new discoveries replace the opportunities which close.

Renewal requires both continuing arrivals of forecasting problems and a positive probability that their interpretation can improve. Better methods can extract more from an existing public record. New businesses, technologies and policies also create questions that existing methods do not fully answer. These features sustain the supply of problems worth investigating; the private reward determines whether investors investigate them.

The paper implements these ideas with a deliberately small model. Section 2 describes the lifecycle of an opportunity and connects the model to the companion papers. Section 3 specifies forecasting problems, interpretation and price formation. Section 4 derives endogenous search and discovery. Section 5 describes the competitive life of a discovered interpretation. Section 6 combines the two sides to obtain the stationary stock result and its large-market and pricing-error corollaries. Section 7 considers limiting cases and comparative statics. Section 8 states the main empirical implications, and Section 9 concludes.

1.1 Contribution and precedents

Lo's (2004) Adaptive Markets Hypothesis emphasises competition, changing environments and the need for investment managers to innovate. Timmermann and Granger (2004) describe efficiency relative to available forecasting models and search technologies, while McLean and Pontiff (2016) document the

erosion of published return predictors. Kandel and Pearson (1995) provide a direct antecedent for differential interpretation of public signals. This paper adds to that literature a model in which the production of superior interpretations is endogenous, and derives the stock of open opportunities from the rate at which they are discovered and the rate at which competition closes them.

Grossman and Stiglitz (1980) show that investors pay for information only when the return to having it covers the cost. Gârleanu and Pedersen (2018) apply this to the market for active management. Managers who pay for information outperform after fees, and investors pay to find them; when finding a manager becomes cheaper, more capital goes to active management, fees fall and prices become more efficient. This paper uses the same condition, that the reward must cover the cost, for a different activity: producing a new interpretation of evidence that is already public. When the reward falls short, nobody searches and no new opportunities appear.

Berk and Green (2004) show how capital flows can absorb investment performance. Loury (1979) supplied the research-race structure. Stein (2008) and Kovbasyuk and Pagano (2022) analyse communication of investment ideas and price discovery. These mechanisms inform the relation between discovery incentives and the subsequent competitive life of an interpretation.

2. The equilibrium process

2.1 The lifecycle of an interpretive opportunity

Figure 1 shows a continuing inflow of forecasting problems. Some contain scope for a superior interpretation, and investors search when discovery is worth its cost. A successful originator establishes a position; subsequent adoption gives the interpretation more influence over price and erodes its remaining alpha. Resolution, obsolescence or declining profitability closes the opportunity. At any date, different episodes occupy different stages of this process.




Figure 1. Continuous renewal of interpretive opportunities

Before discovery, there is latent scope for innovation: a superior interpretation may be producible but has not yet been produced. After successful positioning, there is an open opportunity while the interpretation remains valid and accessible and its expected gross alpha exceeds exploitation costs. The main result of this paper, Theorem 1 in Section 6, concerns the expected stock of open opportunities, i.e. the latter state.

The rate of entry into the stock of open opportunities depends on the incentive to search and the chance of positioning before pre-emption by other investors. The size of the stock also depends on how long each opportunity remains open. The two are linked through the cost of sharing an idea. When that cost falls, the originator shares more widely. This raises her reward, which encourages search, but also means other investors adopt the idea sooner, which shortens how long each opportunity stays open.

2.2 Continuing scope for interpretation

This section recaps the reasons why interpretation of public evidence can keep improving (Arie, 2026a).

Take a single forecasting problem. All investors have access to the publicly available information set (Fama's Φ_t , from 1970). Without other information, an investor can still produce a better forecast either by improving the selection, measurement or weighting of inputs found within the public record, which is almost infinitely large, or by discovering better methods of interpretation (also derived from a universe of possible techniques which is almost infinitely large).

The limit of this process is the perfect forecast of the future produced by Laplace's demon: an intelligence which knows the full present state of the world and can compute every past and future state from it (Laplace, 1814). The demon's interpretive capability is by definition unimprovable, because it is perfect. But real-world forecasts are evidently not close to that standard. Some may claim that there exists another unimprovable limit to what can be understood from the universe of available information and possible forecasting methods but which (a) falls short of perfection, and (b) lies within relevant distance of current practice. That is a claim implicit in efficient markets theory, discussed at length in the companion papers (Arie 2026a, 2026b) and rejected there for lack of evidence.

It follows that scope for better interpretation can come from two sources: further improvement to problems investors are already working on, and new problems which arrive. For simplicity, this model uses only the second. Each episode covers only the first material improvement to a problem, and later improvements to the same problem are left out. They could be added as further episodes, but that would require specifying when they arrive and how they depend on earlier ones. In the model, the supply of opportunities is therefore renewed by new problems alone. (The effect is conservative: if the stock of open opportunities stays above zero with new problems as the only source, it would do so more easily with both.)

New problems can need new interpretive work even as general methods improve. A new technology may change a company's unit costs, or a new policy may change how investment translates into future cash flow. Investors carry forward what they have learned, but it does not fully apply to the new setting. The model assumes that some new problems can still be interpreted better than current general methods allow.

2.3 How the companion-paper results enter the model

Table 1 summarises the companion-paper mechanisms used here.

Table 1. How earlier results enter the equilibrium model

Earlier result	Role in Paper 3
Common public record with heterogeneous interpretation	Allows innovation in the use of public evidence without introducing a private payoff signal.
Capital-weighted price formation	Maps proprietary forecasts into price and makes interpretive influence depend on deployed capital.
Latent dissent	Price does not identify proprietary models; heterogeneous models and priors are permitted.
Alpha-price-improvement identity	Gives economic value to a superior payoff interpretation and links adoption to price improvement.
Research race	Provides the discovery race; Paper 3 adds the marginal search cost c_r and activation threshold.
Dissemination incentive	Lower controlled dissemination cost k can raise the positioned originator's value.
Strategic deployment by one producer	Shows that private deployment can stop short of full price correction when the producer affects price.
Independent deployment by many producers	Converges towards the price-taking adoption limit used as the Paper 3 baseline.
Finite alpha half-life	Gives discovered opportunities finite economic lives even when capability to produce new ideas persists.

3. Market and interpretive environment

3.1 Forecasting problems and interpretations

Forecasting problems arrive according to a stationary process with finite mean rate $\xi \geq 0$. A problem j consists of a public record D_j and an unresolved future payoff X_j . All investors observe D_j . Investor i applies an interpretation χ_{ij} to form forecast $f_{ij} = \chi_{ij}(D_j)$. Interpretation comprises the model, parameters and judgements linking evidence to expected payoff.

With probability $p \in [0,1]$, a new problem contains scope for an economically material improvement. Conditional on that scope, its initial normalized opportunity size α_0 has stationary distribution F . The inflow of potentially improvable problems is ξp . The pair (p, F) describes residual scope relative to

current general capability: p gives its incidence and F its size. Both remain fixed in the stationary baseline.

Searchers know the distribution of potential improvements but learn the size of a particular opportunity only when they discover it. Appendix F states the various maintained assumptions on scope, timing, beliefs and aggregation all in one place.

3.2 Price-taking investors and latent dissent

The baseline investor is small relative to the market. Write $\pi = RP$ for price in payoff units, where P is the quoted price and R the benchmark gross return. A CARA/mean-variance formulation gives forecast-following demand

$$q_i = \lambda_i (f_i - \pi), \quad (3)$$

where λ_i rises with risk tolerance and falls with perceived payoff variance. With a zero-net-supply residual claim, market clearing gives the familiar capital-weighted aggregate

$$\pi = \sum_i w_i f_i, \quad w_i = \lambda_i / \sum_j \lambda_j. \quad (4)$$

An individual price-taker has small λ_i relative to aggregate demand. She can establish a position without materially closing the opportunity on her own. Adoption by other investors subsequently gives the interpretation substantial price weight.

Investors are Bayesian within their own interpretive models; the model does not impose a common prior or common model across investors. Price does not identify another investor's proprietary mapping. The same aggregate can arise from similar forecasts or from offsetting views. Latent dissent (see Arie, 2026a) describes this failure of price to reveal its underlying interpretations.

Investors may learn from price as a noisy aggregate signal, weighting it against their own forecasts according to perceived precision. Proprietary methods nevertheless remain imperfectly revealed.²

3.3 Positioning, pre-emption and diffusion

Before positioning, the price may move to reflect the interpretation through leakage, sufficient independent replication, or other routes into price, preventing the originator from capturing the reward. The pre-emption hazard ζ represents this risk. Replication by a single price-taking investor would not by itself move price; ζ covers replication only to the extent that it helps move the price.

After positioning, adoption can bring forward the first investor's gain and reduce exposure to the unresolved payoff. Under Paper 2's participation and dissemination conditions, she chooses effort m to reach adopters at a running cost of $km^2/2$ per time period.³ Lower k makes this controlled dissemination cheaper. Changes in k do not affect ζ . Lower k means communication is cheaper after she has taken her position; it does not change the risk that she is pre-empted before she has done so.⁴

4. Endogenous search and discovery

4.1 A research race

Consider one problem with scope and $n \geq 1$ investors able to work on it. The set-up follows the research race of Loury (1979), with costs paid per unit of time as in Lee and Wilde (1980) and the companion paper (Arie, 2026b). Investor i chooses a constant research effort $x_i \geq 0$ and discovers the better

² Aumann (1976) shows that agents with a common prior, who know each other's beliefs, cannot agree to disagree. The model avoids this result by allowing investors to use different models to interpret the evidence, so there is no common prior. The fact that price does not reveal other investors' interpretations would not be enough on its own.

³ From Arie (2026b): m is how much effort the investor makes and k is the cost of sustaining that effort.

⁴ A new communication technology may make it cheaper to disseminate an idea after the investor has taken her position (lower k), but also easier for others to copy it before she has done so (higher ζ). The first relationship encourages search; the second discourages it. So the overall effect of such a technology on search, and on how quickly opportunities open and close, can go either way. The model therefore treats the two effects separately.

interpretation at Poisson rate x_i , so her chance of success in any short period is constant. Independent search clocks compete for the first discovery.

Effort costs $C(x_i) = c_r x_i + (k_r/2)x_i^2$ per unit of time, with $c_r > 0$ and $k_r > 0$. This functional form is chosen for simplicity. The quadratic term, as in the companion paper, makes each extra unit of effort more costly than the last. The linear term is new: it means even the first unit of effort has a cost, so search can be unprofitable. The first discovery ends the race: the other searchers learn that the problem has been solved, though not the interpretation, and stop searching.⁵

Let μ_0 be the rate of resolution before discovery and r the discount rate. After discovery, positioning occurs at rate ρ and pre-emption at rate ζ . Write $W_p(a; k)$ for the private value once positioned with initial opportunity size a , and assuming she has optimised dissemination effort m . Before discovery, the expected continuation value is $W_p = E_F[W_p(\alpha_0; k)]$, and the expected discovery prize is

$$V = E_F[V(\alpha_0)] = [\rho / (r + \rho + \zeta)] W_p. \quad (5)$$

V is what the new interpretation is worth to the originator at the moment she finds it, before she has taken her position. V is smaller than W_p firstly because the investor may be pre-empted before she positions, and secondly because positioning takes time so the value is discounted. The fraction is therefore a discounted probability of positioning first: without discounting ($r = 0$) the fraction would be $\rho/(\rho + \zeta)$, i.e. the plain chance that she positions before she is pre-empted. Adding $r > 0$ reduces the expectation further simply because the reward comes later.

Search therefore depends on the expected prize across possible discoveries. A valuable interpretation may be producible without being worth the cost of finding it.

⁵ If rivals did not learn of a discovery and kept searching, no one would search when $V \leq c_r$, and zero search would not be an equilibrium when $V > c_r$; equilibrium effort would then differ from (7).

Let $\tilde{V} = V - c_r$ and X_{-i} be other investors' total effort.⁶ If investor i keeps her effort at x_i until the search ends, the expected value of the search is:

$$J_i = [x_i \tilde{V} - (k_r/2)x_i^2] / [b + X_{-i} + x_i], \quad b = r + \mu_0. \quad (6)$$

The numerator is the expected prize per unit of time, net of the cost of effort. The denominator converts this flow into a present value: the search ends when anyone discovers the answer or the problem is resolved, and future payoffs are discounted at rate r . The positive marginal cost c_r allows search to be inactive even when discovery has positive private value.

4.2 Equilibrium search

If $V \leq c_r$, every investor's unique best response is zero search. If $V > c_r$, the unique symmetric stationary equilibrium has $x_i = x^* > 0$, where

$$k_r(n - 1/2)x^{*2} + [k_r b - (n - 1)\tilde{V}]x^* - b\tilde{V} = 0. \quad (7)$$

For $\tilde{V} > 0$, the constant term is negative and the leading coefficient positive, giving exactly one positive root. Search rises with V and with the number of searchers n , and falls with c_r or k_r . At $V = c_r$, effort reaches zero continuously. Appendix A derives both regimes and the positive root.

Discovery precedes resolution with probability $nx^*/(\mu_0 + nx^*)$. Conditional on discovery, positioning precedes pre-emption with probability $\rho/(\rho + \zeta)$. Potentially improvable problems arrive at rate ξp , so successful entries into the economically open state occur at rate

$$v^* = \xi p \cdot [nx^* / (\mu_0 + nx^*)] \cdot [\rho / (\rho + \zeta)]. \quad (8)$$

⁶ Effort is measured by its result: one unit of effort is whatever produces a discovery rate of one per unit of time. So $x = 0.5$ means enough work to find the answer at a rate of 0.5 per year, or after two years on average, whatever that takes in analysts, data and time.

The distribution of opportunity sizes affects entry through the expected prize and equilibrium search effort. Searchers choose effort before they know the size of the discovery.

If $V \leq c_r$, $v^* = 0$ even when $\xi p > 0$. If $V > c_r$, $\xi p > 0$ and positioning is possible, $v^* > 0$. A reduction in k can raise W_p and hence entry under the dissemination conditions below. Greater ζ reduces both the prize and the chance of successful positioning.

In sum, investors search only when the expected prize from discovery is greater than the cost of the first unit of effort. When it is, there is a single equilibrium level of search, which rises with the prize and falls with the cost of searching, and which declines gradually to zero as the prize falls towards that cost.

Equation (8) converts this search effort into the rate at which new opportunities open. That rate is positive only if problems with scope keep arriving, investors find the better interpretation before the problem is resolved, and the discoverer takes her position before the price moves.

5. The competitive life of a discovered interpretation

5.1 From one producer to many price-takers

Once the originator has taken her position, the better interpretation becomes an open alpha opportunity.

She is a price-taker: her own buying has almost no effect on price, so nearly all of the initial gap between her forecast and the price is still there after she has taken her position.

The companion paper (Arie, 2026b) had considered a producer large enough to move price. She holds back, taking a smaller position than would make price most accurate, because further buying would reduce her return on the position she already holds. Several independent producers hold back less, and as their number grows, their combined position approaches the level that makes price most accurate. The model here moves to the more common case in which investors, including the originator, are price-takers, so none holds back. The one difference from the companion paper's many-producer case is that adoption

stops once the remaining alpha no longer covers the exploitation cost c_e , so price stops short of the most accurate level.

Adopters do not need to reproduce all of the originator's reasoning. They may be given the method, copy a trading rule, or infer enough from a disclosed position to revise their own forecast. Each adoption brings capital that gives the interpretation more weight in price, moving price part of the way towards the better forecast: in the model, each adoption leaves a fraction η of the remaining alpha.

5.2 Half-life and economic closure

Let $\alpha(u; a)$ be remaining normalized gross alpha at age u after positioning, conditional on initial size $\alpha_0 = a$. Adoption (of the idea by other investors) and resolution (by the arrival of confirmatory events) reduce it in expectation, although individual paths can jump because both types of event exhibit randomness.

An opportunity remains economically open while its gross alpha exceeds exploitation costs and its interpretation remains valid and accessible. Let $\tau(a)$ be the age at which the opportunity closes. From then on, its remaining alpha $\alpha(u; a)$ is counted as zero, since whatever is left no longer covers the cost of trading on it.

Average over both size and the conditional adoption and resolution paths. The mean life is $T = E_F[E[\tau(\alpha_0) | \alpha_0]]$, and the expected lifetime normalized opportunity is

$$Q = E_F[E[\int_0^{\tau(\alpha_0)} \alpha(u; \alpha_0) du | \alpha_0]]. \quad (9)$$

Use the chosen contact rate $m = m^*(a; k)$ from Paper 2 as the arrival rate of adopters. Each adoption leaves a fraction η of remaining alpha, and resolution arrives at rate μ . Before applying the exploitation cutoff (i.e. the cutoff c_e , where remaining alpha no longer covers the exploitation cost) the conditional mean path is derived in Appendix D and results as follows:

$$E[\alpha(u) | \alpha_0 = a] = a \exp\{-[\mu + m(1 - \eta)]u\}. \quad (10)$$

Write $\delta(a) = \mu + m^*(a; k)(1 - \eta)$. The conditional mean has half-life $\ln(2)/\delta(a)$. Different initial sizes can induce different dissemination choices, so their average need not have a single exponential half-life.

Appendix D derives T and Q for the process stopped at economic closure.

5.3 Diffusion and the value of discovery

The originator chooses dissemination to bring forward the gain on her position. Where dissemination is worthwhile, a lower cost k raises both her chosen contact rate m^* and her continuation value W_p . Where she chooses no dissemination, changing its cost has no effect. Appendix B derives these results from Paper 2's timing problem.

The link to the opportunity's lifetime requires one additional assumption: the channel established by the originator continues to generate adoption at her chosen rate after she sells. Follow-on adoption is not separately optimised. Under this persistence assumption, the same m^* determines the private prize, T and Q. Cheaper dissemination raises the prize and shortens the opportunity's life; where search pays, it also brings more discoveries into the market.

The persistence assumption makes market-wide erosion depend only on the originator's dissemination choice. Adopters do not choose to spread the interpretation further, and no second searcher looks for it once the originator is positioned. Both would raise the rate of adoption. The resulting T and Q are therefore upper bounds on the life and lifetime opportunity of a discovered interpretation. For

6. Stationary competitive equilibrium

6.1 The stock of open opportunities

In a stationary market, discoveries enter and opportunities close at equal average rates. The number open at any date depends on the rate of entry and how long each remains open.

Each positioned discovery remains economically open for its random lifetime. Little's law gives the expected count

$$N^* = v^* T, \quad T = E_F[E[\tau \mid \alpha_0]]. \quad (11)$$

Equation (8) determines whether the entry rate is positive. Under $\xi p > 0$, $V > c_r$ and the maintained timing assumptions, $v^* > 0$ and therefore $N^* > 0$. In the no-search regime, the stationary count is zero.

Summing the normalized gross alpha remaining across open opportunities and applying the stationary reward identity gives

$$O^* = v^* Q. \quad (12)$$

O^* sums the normalized alpha remaining across open opportunities. Positive entry and positive lifetime opportunity imply a positive stationary stock.

6.2 Theorem and interpretation

Theorem 1 (Stationary search and open opportunities). Under the maintained assumptions in Appendix F, the unique symmetric search equilibrium has $x^* = 0$ if $V \leq c_r$ and otherwise the positive root of (7). Successful entry is given by (8), and the stationary stocks are $N^* = v^* T$ and $O^* = v^* Q$. Both stocks are positive exactly when $\xi p > 0$ and $V > c_r$. They are zero if $\xi p = 0$ or $V \leq c_r$.

Appendix A proves the search threshold and the equilibrium level of effort, and Appendix B derives the expected prize. Equation (8) follows because discovery, resolution, positioning and pre-emption are independent random events, and the stocks follow from the steady-state relationships in Appendix C. The threshold matters economically: an interpretation that would be useful to the market can go undiscovered if the reward to the investor who finds it does not cover the cost of the first unit of search. When search is active, competition closes each opportunity over time, and search opens new ones to replace them.

Adoption spreads the use of interpretations that have already been found. Research adds new ones. The two are connected through W_p , the value of an opportunity to the investor who finds it: how quickly others adopt her interpretation affects that value, and so affects the incentive to search for the next one.

The total stock of open opportunities can therefore stay steady on average while individual opportunities keep opening and closing.

Note: The result could be read as assuming its conclusion, since the model assumes that problems with scope keep arriving and concludes that some opportunities are open. But the assumption is necessary, not sufficient. With the same inflow of scope, the stock is zero when the reward to discovery does not cover the cost of the first unit of search, so whether scope becomes open opportunities is an equilibrium outcome. Competition among searchers also works against the usual intuition. The harder an investor searches, the more each unit of effort costs, so without time pressure she would search slowly. Each additional rival raises the chance that someone else finds the answer first, so each searcher searches harder, and more searchers find opportunities faster. The life of an opportunity does not depend on the number of searchers, so the stock of open opportunities rises with the number of competitors. This is a known property of research races in which costs are paid while the race lasts (Lee and Wilde, 1980; Appendix A). For the stationary expected number of open opportunities to be zero, at least one of the following must hold: scope for better interpretation has run out, new problems have stopped arriving, the reward to discovery does not cover the cost of search, research is prohibitively costly, or opportunities close as soon as they open. Continuing scope is itself an assumption about the world, argued in Section 2.2 and in Arie (2026a); the paper does not derive it.

6.3 A large-market corollary

Theorem 1 is about the average number of open opportunities. It does not say that a market always has at least one open at any given moment. A stronger statement is possible for a large market, if opportunities in different areas of forecasting are independent of one another, or nearly so. If opportunities open at random times, as a Poisson process, and their lifetimes are independent, then the number open at any moment has a Poisson distribution with mean N^* . The chance that none is open is then

$$\Pr(N = 0) = \exp(-N^*). \quad (13)$$

This chance falls quickly as N^* rises. With $N^* = 1$ it is about 37%; with $N^* = 5$ it is under 1%. So in a market with many independent areas of forecasting, and therefore a large N^* , it is very unlikely that every opportunity is closed at the same time.

Correlation weakens this. A single macroeconomic question or a technology change affecting a whole sector can create many related trades, which behave more like one opportunity than many. Correlation therefore reduces the effective number of independent opportunities. Equation (13) is a benchmark for the independent case; the paper does not claim that opportunities on individual securities are independent.

6.4 Price accuracy as a conditional corollary

Under one special case, alpha has a direct meaning for price accuracy. The companion paper (Arie, 2026b) shows that when the better forecast is correctly calibrated ($A = B$ in that paper's notation), the alpha remaining in an opportunity equals exactly how much more accurate price would become if it fully reflected the better forecast, measured by the reduction in squared forecast error. If open opportunities don't interact and are measured in the same units, O^* is then the expected total improvement in price accuracy that could still be gained from interpretations currently open.

This link holds only in that special case. Theorem 1 does not need it. And alpha does not always come from improving accuracy: an investor can profit from others adopting an interpretation even if it is wrong, by selling to them at a higher price. Appendix E sets out the equality and its units.

7. Limiting cases and comparative statics

7.1 When the opportunity stock disappears

A problem has interpretive scope if some interpretation of its public record would forecast the payoff better than the price does, by enough that the resulting alpha exceeds the cost of exploiting it ($\alpha_0 > c_e$). It exists before anyone has found that interpretation; searching is the attempt to find it. p is the share of new problems that have it.

Interpretive scope defined this way can disappear. If $p = 0$, new problems offer no material improvement and $v^* = N^* = O^* = 0$. Holding other primitives and moments fixed, the same stocks tend to zero as p tends to zero. This is the interpretive-completeness limit discussed in Section 2.2, of which Laplace's demon is the extreme case.

Renewal can cease. If $\xi = 0$, no new episodes enter and the stationary stock is zero. An inherited finite stock of unresolved problems or open opportunities can persist during transition, but $\mu_0 > 0$ and finite opportunity lives prevent it from sustaining the long-run stock.

Search can be uneconomic despite positive scope. If $V \leq c_r$, all searchers choose zero effort and no new opportunities are opened. This efficient-market-style outcome is an equilibrium within the model: private returns to finding a superior interpretation do not cover the marginal cost of looking. Higher ζ can produce this outcome at a finite hazard, since $V = \rho W_p / (r + \rho + \zeta)$. With W_p fixed, search is active only if $\zeta < \rho W_p / c_r - r - \rho$. If the right-hand side is zero or negative, no one searches even when there is no risk of pre-emption: the reward is too small for other reasons.

Economic lives can also approach zero. With finite ξp , the bound $v^* \leq \xi p$ means that $T \rightarrow 0$ drives $N^* \rightarrow 0$, and $Q \rightarrow 0$ drives $O^* \rightarrow 0$. Instantaneous absorption after positioning may leave the originator a reward even as the open stock vanishes. Absorption before positioning reduces that reward and can stop search.

Finally, x^* and v^* tend to zero as k_r tends to infinity, holding other primitives fixed. A zero stationary expected stock can therefore arise through exhausted scope, absent renewal, insufficient private reward, prohibitive research costs or vanishing economic lives. The no-innovation region does not imply that every latent interpretation is already reflected in price; it says that none is worth producing privately within this model.

7.2 Competition and turnover

Consider a reduction in the controlled dissemination cost k , holding ζ , ξ , p and F fixed. Under Paper 2's adoption conditions and $Z > 0$, lower k raises optimal m^* and W_p . If the expected V exceeds c_r , search and the rate of successfully positioned discoveries rise. If V remains below the threshold, cheaper dissemination has no effect on entry; crossing the threshold activates search.

Under the dissemination-persistence assumption, the higher chosen adoption rate reduces T and Q . Lower k therefore raises the inflow while shortening opportunity lives. The expected normalized stock is

$$O^*(k) = v^*(k) Q(k), \quad \zeta \text{ held fixed.} \quad (14)$$

The sign of the stock effect is ambiguous. Where lower k raises the expected prize and search is active, it increases the flow of entries, v^* . In stationarity, the same number closes per unit time on average. Shorter lives also raise turnover relative to the count of open opportunities, $v^*/N^* = 1/T$.

Higher pre-emption ζ works the other way. It lowers the prize V and the chance that the discoverer takes her position in time, so fewer new opportunities open while search is active, and search may stop altogether. A technology which lowers k but raises ζ therefore has an effect that can go either way. Ways of communicating that the originator controls after she has taken her position should increase the incentive to search, provided the conditions stated above hold; anything that makes it easier for others to copy the idea before she has positioned should reduce it.

7.3 General capability and problem specificity

General forecasting capability may improve over time. If its benefits transferred perfectly to future problems of unchanged difficulty, p or the distribution of α_0 could shrink. The expected prize might then fall below c_r before all scope was exhausted.

The stationary benchmark holds ξ , p and F fixed relative to current capability. It describes settings in which novel questions continue to require specific interpretive work. A positive stationary stock requires

continuing arrivals with positive probability of scope, an expected prize above the marginal cost of starting search, and positive finite opportunity lives. Further claims about technological progress would require dynamics for these primitives.

8. Empirical implications

The model suggests that individual opportunities should lose alpha as their methods spread and capital adopts them. Their decay need not imply a declining market-wide stock, since new discoveries can replace closed opportunities. Successful interpretive innovations should sometimes improve forecasts using evidence that was public before the method was discovered.

The dissemination distinction adds a testable qualification. Cheaper communication controlled by a positioned originator should raise discovery incentives when it raises the expected prize above the search threshold. Easier copying before positioning should reduce those incentives. Faster turnover is predicted for the former change under active search; generic measures of communication technology may combine opposing effects.

Testing these implications requires dated measures of methods, forecasts, adoption and implementation costs. The key comparison is between the erosion of existing opportunities and the arrival of their successors. The model also changes what an observation of open opportunities indicates. In this equilibrium, a positive stock of them is produced by competitive search, so finding profitable interpretations in a market is consistent with competition working rather than evidence that it has failed. The stationary expected stock of open opportunities is zero when scope for improvement has gone, new problems have stopped arriving, or the reward to discovery does not cover the cost of searching.

9. Conclusion

Competition erodes the value of a discovered interpretation. The prospect of earning that value also makes investors search for better interpretations. Controlled dissemination connects the two: it brings forward the originator's reward while accelerating the adoption that closes the opportunity.

A market with a zero stationary expected stock of open opportunities is a corner of this model. It arises when interpretation offers no further scope, new forecasting problems cease to arrive, or the reward to discovery fails to cover the marginal cost of search. Instantaneous absorption would be a further limiting case. But competition alone does not force the market into any of these states.

Where new problems continue to offer scope and search pays, competition closes particular opportunities while opening others. This is the central claim of the trilogy: the use of public evidence can improve, the reward to each improvement can be temporary, and yet the most effective investors can still earn. This gives the paradoxical conclusion that the most important signature of competition is a positive stock of open alpha opportunities.

Appendix A. Search equilibrium

Consider $n \geq 1$ originators searching one live problem with positive scope. Effort $x_i \geq 0$ gives discovery rate x_i and flow cost $c_r x_i + k_r x_i^2/2$. Set $b = r + \mu_0 > 0$ and $\tilde{V} = V - c_r$, where V is the expected prize over F . The other searchers' aggregate rate is X_{-i} .

For constant effort the stationary expected value is

$$J_i(x_i; X_{-i}) = [x_i \tilde{V} - (k_r/2)x_i^2] / [b + X_{-i} + x_i]. \quad (A1)$$

Write $d = b + X_{-i} > 0$. The derivative of J_i has the sign of $h(x_i) = \tilde{V}d - k_r dx_i - (k_r/2)x_i^2$. If $\tilde{V} \leq 0$, every positive effort has negative value and the unique optimum is zero. If $\tilde{V} > 0$, $h(0) > 0$ and h is strictly decreasing to a negative value. There is one positive best response.

Under symmetry, $X_{-i} = (n - 1)x$. The first-order condition is

$$k_r(n - 1/2)x^2 + [k_r b - (n - 1)\tilde{V}]x - b\tilde{V} = 0. \quad (A2)$$

For $\tilde{V} > 0$, the leading coefficient is positive and the constant term negative. Hence there is one positive root. Let $a_r = k_r(n - 1/2)$ and $d_r = k_r b - (n - 1)\tilde{V}$. It is $x^* = [\sqrt{(d_r^2 + 4a_r b \tilde{V})} - d_r]/(2a_r)$. At $\tilde{V} \leq 0$ the constrained solution is $x^* = 0$; the positive-root formula is used only in the active regime.

The polynomial derivative at its positive root is $\sqrt{(d_r^2 + 4a_r b \tilde{V})} > 0$. Its derivative with respect to $\tilde{V}$ is $-(n - 1)x^* + b] < 0$, and with respect to k_r is $(n - 1/2)x^{*2} + b x^* > 0$. Implicit differentiation therefore gives $\partial x^*/\partial V > 0$, $\partial x^*/\partial c_r < 0$ and $\partial x^*/\partial k_r < 0$. The root approaches zero as $V \downarrow c_r$. Independent search and resolution clocks give discovery probability $n x^*/(\mu_0 + n x^*)$, including zero at $x^* = 0$.

Treating n as continuous, the best-response condition gives $\tilde{V} > k_r x^*$, so the derivative of (A2) with respect to n , $x^*(k_r x^* - \tilde{V})$, is negative at the positive root. Hence $\partial x^*/\partial n = x^*(\tilde{V} - k_r x^*)/\sqrt{(d_r^2 + 4a_r b \tilde{V})} > 0$: individual and total search rise with n , and so do the discovery probability and v^* . This is the result of

Lee and Wilde (1980) for research races in which costs are paid while the race lasts. In Loury's (1979) version, where each firm pays its research cost at the start, individual effort falls as the number of firms rises.

No fixed participation cost is needed for the inactive regime.

Appendix B. Positioning, pre-emption and the discovery prize

Conditional on $\alpha_0 = a$, let positioning and pre-emption occur at independent rates ρ and ζ . Pre-emption closes this episode and pays zero; successful positioning delivers $W_p(a; k)$. Size is independent of these clocks. Averaging the discounted positioning reward over F gives

$$V = E_F \left[\int_0^\infty \rho e^{-(r+\rho+\zeta)t} W_p(\alpha_0; k) dt \right]. \quad (B1)$$

For fixed W_p , $\partial V / \partial \zeta = -\rho W_p / (r + \rho + \zeta)^2 < 0$ when $W_p > 0$. Faster positioning weakly raises V , strictly if $W_p > 0$ and $r + \zeta > 0$. Physical entry uses $\rho / (\rho + \zeta)$, whereas the private prize also discounts delay. Combining physical positioning success with the research success probability and inflow ξp gives equation (8).

For controlled dissemination, use Paper 2's fixed-position timing problem conditional on a . Let $S(a)$ be the utility gain from selling after adoption, $H(a)$ the expected utility gain at resolution, and μ the resolution hazard after positioning. A qualified recipient must satisfy the participation condition. Hold its mandate, the originator's position and these payoffs fixed as k varies. Write $b_D = r + \mu$ and $Z(a) = rS(a) + \mu[S(a) - H(a)]$.

$$b_D W_p(a; k) = \mu H(a) + \max_{m \geq 0} \{m[S(a) - W_p(a; k)] - km^2/2\}. \quad (B2)$$

For $Z(a) > 0$, Paper 2's equation (E2) gives $m^*(a; k) = \sqrt{(b_D^2 + 2Z(a)/k) - b_D}$. For $Z(a) \leq 0$, $m^* = 0$ and $W_p(a; k) = \mu H(a) / b_D$. With $Z > 0$, differentiating the solution for effort and applying the envelope argument to (B2) gives

$$\partial m^*/\partial k = -Z / [k^2 \sqrt{(b_D^2 + 2Z/k)}] < 0, \quad (B3)$$

$$\partial W_p(a; k)/\partial k = -m^{*2} / [2(b_D + m^*)] < 0. \quad (B4)$$

Both derivatives are zero in the inactive dissemination region. Averaging (B4) over F gives the expected-value result where differentiation is valid. With ζ fixed, $\partial V/\partial k = [\rho/(r + \rho + \zeta)] \partial W_p/\partial k$. If $V > c_r$, Appendix A then implies that lower k raises x^* and v^* . This argument holds private payoffs and recipient participation fixed; it makes no claim about a technology that also raises pre-emption.

Appendix C. Stationary stocks

Let successfully positioned discoveries enter in a stationary process with mean rate v^* . Entry sizes have distribution F . Each size a has conditional lifetime $\tau(a)$ and remaining normalized alpha $\alpha(u; a)$, zero after closure. Assume marks are independent of entry times and their means T and Q are finite. Common normalization and separability permit addition across opportunities.

For entry age u , the probability of remaining open is $E_F[\Pr(\tau > u \mid \alpha_0)]$. Integrating it over all ages gives T . Thus the stationary expected count is

$$N^* = v^* \int_0^\infty E_F[\Pr(\tau > u \mid \alpha_0)] du = v^* T. \quad (C1)$$

The same argument integrates each entry's expected remaining alpha over its possible ages. With $Q = E_F[E[\int_0^\tau \alpha(u; \alpha_0) du \mid \alpha_0]]$, the normalized opportunity stock is

$$O^* = v^* Q. \quad (C2)$$

The identities hold in both search regimes. If $\xi_p = 0$ or $V \leq c_r$, then $v^* = 0$ and both stocks vanish. Under $\xi_p > 0$ and $V > c_r$, entry is positive and finite; positive finite T and Q then give positive finite stocks. The distribution F enters both the incentive to create opportunities, through V , and their expected duration and size, through T and Q .

If entries are Poisson and lifetimes are independent, the system is an $M/G/\infty$ queue. The stationary number open is Poisson with mean v^*T , yielding $\Pr(N = 0) = \exp(-N^*)$. Common shocks and correlated opportunity lives invalidate the exact Poisson statement, which is why the main text treats it as a large-market benchmark.

Appendix D. Adoption, economic closure and dissemination costs

Condition on $\alpha_0 = a$ and set $m = m^*(a; k)$, the optimal contact rate in Appendix B. The first adoption or resolution ends the originator's holding and determines her private payoff. Under the persistence assumption, subsequent adoptions continue at the same rate m without further effort or cost from her. Resolution occurs at the same rate μ throughout. Each adoption leaves fraction $\eta \in [0, 1)$ of remaining alpha. Before the exploitation cutoff, the unresolved path is $a\eta^{N_m(u)}$, where $N_m(u)$ is Poisson with mean $m \cdot u$ and $\eta^0 = 1$. Then

$$E[\eta^{N_m(u)}] = \exp\{-m(1 - \eta)u\}. \quad (D1)$$

Including survival to factual resolution and writing $\delta = \mu + m(1 - \eta)$ gives

$$E[\alpha(u) \mid \alpha_0 = a] = a \exp(-\delta u). \quad (D2)$$

For common m , μ and η , averaging over F gives $E[\alpha(u)] = E_F[\alpha_0]e^{-\delta u}$. If decay rates depend on a , the corresponding expression is $E_F[\alpha_0 e^{-\delta(a)u}]$. The conditional mean-path half-life is $\ln(2)/\delta(a)$. Its crossing of c_e , for $a > c_e > 0$, occurs at $\ln(a/c_e)/\delta(a)$. This is not the realised lifetime.

Without an exploitation cutoff, $Q = E_F[\alpha_0/\delta(\alpha_0)]$. With a positive cutoff c_e , both the opportunity's lifetime and its remaining alpha must be stopped at closure. The following expressions use the chosen $m^*(a; k)$, not an independent exogenous adoption rate.

For a positive cutoff c_e and $a > c_e$, let $K(a)$ be the smallest positive integer for which $a\eta^{K(a)} \leq c_e$, with $K(a) = 1$ when $\eta = 0$. The open states have adoption counts $j = 0, \dots, K(a) - 1$. At each state, expected holding

time before adoption or resolution is $1/(\mu + m)$, and the probability of reaching state j is $[m/(\mu + m)]^j$.

Hence

$$T(a) = \sum_{j=0}^{K(a)-1} [m/(\mu + m)]^j / (\mu + m), \quad (D3)$$

$$Q(a) = a \sum_{j=0}^{K(a)-1} [m\eta/(\mu + m)]^j / (\mu + m). \quad (D4)$$

The unconditional quantities are $T = E_F[T(\alpha_0)]$ and $Q = E_F[Q(\alpha_0)]$. Thus W_p , T and Q depend on the same size distribution and chosen contact rates. The repeated-adoption process extends Paper 2's first-recipient timing problem through the persistence assumption; it does not derive continuing dissemination effort after the originator exits.

For fixed a , couple two adoption processes using the same resolution time and add adoption events to obtain the higher rate. Remaining alpha and closure time are then no greater on any path. More explicitly, with $h = m/(\mu + m)$, (D3) gives $T(a) = (1 - h^{K(a)})/\mu$ and $\partial T(a)/\partial m = -K(a)h^{K(a)-1}/(\mu + m)^2$. For $m > 0$ this is strictly negative. Equation (D4) is also strictly decreasing in m : an additional adoption lowers remaining alpha or closes the opportunity earlier.

Since $\partial m^*/\partial k < 0$ wherever $Z(a) > 0$, $T(a)$ and $Q(a)$ rise with k in that region and are unchanged where dissemination is inactive. Averaging gives $\partial T/\partial k > 0$ and $\partial Q/\partial k > 0$ when a positive F-mass has active dissemination and differentiation under the expectation is valid. Equivalently, lower k shortens mean life and reduces lifetime opportunity. These results hold F , μ , η and c_e fixed. Bounds $T \leq 1/\mu$ and $Q \leq E_F[\alpha_0]/\mu$ ensure finite means under Appendix F.

Appendix E. Pricing-error corollary

For a second-order payoff interpretation, separate the focal forecast f_s from the background aggregate f_0 and let $d = f_s - f_0$. Write $A = E[d(X - f_0)]$ and $B = E[d^2]$. If the interpretation has price weight θ , the

market forecast is $f_0 + \theta d$. In the calibration $A = B$, the focal interpretation is aligned with the background forecast error in the sense used in Paper 2.

Let $G = 1 - \theta$ be the remaining price-weight gap and $L(f) = E[(X - f)^2]$. Paper 2 gives

$$\alpha = BG^2 = L(\pi) - L(f_s). \quad (E1)$$

Both sides have squared-payoff units: α is profit divided by demand intensity. In common payoff units, summing over separable open claims gives equality between their normalized opportunity sum and the corresponding correctable forecast-loss sum. Taking stationary expectations gives $O^* = v^*Q$. This neither measures dollar profit nor includes undiscovered scope. Without $A = B$, or with material interactions among claims, the equality need not hold; the stock identities remain valid under their own assumptions.

Appendix F. Maintained assumptions

The baseline uses the following assumptions throughout. They separate the economic mechanism from the conditions needed to aggregate it.

1. Scope and size. Problems arrive stationarily at finite rate $\xi \geq 0$. A share $p \in [0,1]$ has identifiable scope for improvement. Searchers know this scope and the stationary conditional distribution F , but learn α_0 only at discovery. F is supported on $\alpha_0 > c_e > 0$ and has a finite mean. Recognising scope does not reveal the superior interpretation. Each episode models one discovery; repeated improvements to a problem would require additional episodes and their dependence to be specified. Search effort is directed only at problems identified as having scope, so unproductive search is outside the model; allowing for it would lower the expected return to research effort.

2. Timing and search. The number of searchers n is a positive integer; $\mu_0, \rho, \mu, c_r, k_r$ and k are positive and finite; r and ζ are non-negative and finite. Discovery, positioning and pre-emption clocks are independent of size and of one another. Resolution before discovery occurs at μ_0 . During positioning, ζ includes every

event that closes the episode without a prize; no additional μ_0 clock applies. Pre-emption is absorbing, so partial leakage requires an additional state. Successful entry retains distribution F . The first discovery ends the race: the other searchers learn that the problem has been solved, though not the interpretation, and stop searching.

3. Interpretation and learning. Investors use Bayesian reasoning within their own models, without a common-prior or common-model restriction across investors. Price learning leaves proprietary methods imperfectly revealed. Non-invertibility of price alone would not establish persistent disagreement under common-prior assumptions.

4. The private prize. Conditional on size, the originator's initial position, qualified recipient's mandate and participation, sale payoff $S(a)$, and resolution payoff $H(a)$ are fixed as k varies. The originator exits at the first adoption or resolution. Expected private values are finite and non-negative. Where derivative statements are averaged over F , differentiation under the expectation is assumed valid.

5. Dissemination persistence. Conditional on size, qualified adoptions arrive at the chosen rate $m^*(a; k)$. The communication channel continues at this rate after the originator sells, without further effort or cost from her. This is an assumption linking Paper 2's first-recipient prize to subsequent adoption, not a further optimal-control result. Adoption and resolution are independent; μ is constant, and each adoption leaves the fixed fraction $\eta \in [0,1)$ of remaining α . Economic closure occurs at resolution or when remaining α first falls to c_e . Comparisons in k hold ζ and the other primitives fixed.

6. Aggregation. Opportunities are separable: exploiting one does not mechanically alter another's α or private value. Search has no shared capacity constraint across episodes. Conditional lifetime paths are stationary and independent of entry times. The exact Poisson corollary additionally requires Poisson entry and independent lifetimes.

7. Units. Claims share a payoff numeraire and demand-intensity normalization. Normalized alpha is gross expected trading profit per unit of demand intensity, with squared-payoff units in Paper 2; it is neither a return percentage nor dollar profit. T has time units, Q has normalized-alpha-times-time units, and O^* sums normalized alpha across opportunities. Private prizes and costs use Paper 2's utility units: c_r is a private value, k_r a private-value-times-time coefficient, and $C(x)$ a flow cost. The normalized stock is not the private prize.

References

- Arie, Sam, 2026a, Interpretation: How Price Formation Works When the Same Public Evidence Supports More Than One Reading, working paper. Available at SSRN <https://ssrn.com/abstract=7477019>
- Arie, Sam, 2026b, Competition: A Theory of Markets in Which Alpha Is a Feature Rather Than a Bug, working paper. Available at SSRN: <https://ssrn.com/abstract=7477021>
- Aumann, Robert J., 1976, Agreeing to Disagree, *Annals of Statistics* 4(6), 1236–1239.
- Berk, Jonathan B., and Richard C. Green, 2004, Mutual Fund Flows and Performance in Rational Markets, *Journal of Political Economy* 112(6), 1269–1295.
- Fama, Eugene F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance* 25(2), 383–417.
- Gârleanu, Nicolae, and Lasse Heje Pedersen, 2018, Efficiently Inefficient Markets for Assets and Asset Management, *Journal of Finance* 73(4), 1663–1712.
- Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the Impossibility of Informationally Efficient Markets, *American Economic Review* 70(3), 393–408.
- Kandel, Eugene, and Neil D. Pearson, 1995, Differential Interpretation of Public Signals and Trade in Speculative Markets, *Journal of Political Economy* 103(4), 831–872.
- Kovbasyuk, Sergey, and Marco Pagano, 2022, Advertising Arbitrage, *Review of Finance* 26(4), 799–827.
- Laplace, Pierre-Simon, 1814, *Essai philosophique sur les probabilités* (Courcier, Paris).
- Lee, Tom, and Louis L. Wilde, 1980, Market Structure and Innovation: A Reformulation, *Quarterly Journal of Economics* 94(2), 429–436.

Little, John D. C., 1961, A Proof for the Queuing Formula $L = \lambda W$, *Operations Research* 9(3), 383–387.

Lo, Andrew W., 2004, The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective, *Journal of Portfolio Management* 30(5), 15–29.

Loury, Glenn C., 1979, Market Structure and Innovation, *Quarterly Journal of Economics* 93(3), 395–410.

McLean, R. David, and Jeffrey Pontiff, 2016, Does Academic Research Destroy Stock Return Predictability? *Journal of Finance* 71(1), 5–32.

Stein, Jeremy C., 2008, Conversations among Competitors, *American Economic Review* 98(5), 2150–2162.

Timmermann, Allan, and Clive W. J. Granger, 2004, Efficient Market Hypothesis and Forecasting, *International Journal of Forecasting* 20(1), 15–27.

Statement on the use of AI

This paper is the third in a set of working papers developed from the author's earlier essay on how financial markets work. The first paper makes interpretation of the public record explicit in price formation; the second develops the competitive production, deployment and erosion of interpretations; the present paper places those mechanisms in a stationary equilibrium.

OpenAI's ChatGPT and Anthropic's Claude were used for literature searching, analytical and methodological reasoning, challenging proposed hypotheses, formulating and checking the model, developing derivations, drafting, revision and critical review. Each tool was frequently asked to critique the other's work as well as the author's.

The author set the research objectives and overarching argument, assessed and resolved the AI input, drafted, revised and line-edited the text, checked references against sources and decided the goals, content, style, structure and conclusions. The proofs and derivations have not yet been independently checked except by the author. The author takes responsibility for the content, conclusions and any remaining errors. The paper is published in working form to invite review and discussion.