International Journal of Applied and Advanced Scientific Research (IJAASR)
International Peer Reviewed - Refereed Research Journal, Website: www.dvpublication.com
Impact Factor: 5.655, ISSN (Online): 2456 - 3080, Volume 10, Issue 2, July - December, 2025

PRICE DEVIATION AND PUNISHMENT CHARACTERISTICS ASSOCIATED WITH ALGORITHMIC COLLUSION STABILITY IN SIMULATED DUOPOLY PRICING MARKETS

Divya Vishwanath* & Paul Johnson Asamoah**
* Janardhan Rai Nagar Rajasthan Vidyapeeth, Udaipur, Rajasthan, India
** Power Life University College, Accra, Ghana, West Africa

Cite This Article: Divya Vishwanath & Paul Johnson Asamoah, "Price Deviation and Punishment Characteristics Associated with Algorithmic Collusion Stability in Simulated Duopoly Pricing Markets", International Journal of Applied and Advanced Scientific Research, Volume 10, Issue 2, July - December, Page Number 149-157, 2025.
Copy Right: © DV Publication, 2025 (All Rights Reserved). This is an Open Access Article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.
DOI: https://doi.org/10.5281/zenodo.24039545

Abstract:

This analysis examines whether resistance to price deviations supports consistent interpretations of algorithmic coordination in the international competition policy debate. A descriptive reanalysis uses 95 published scenario summaries from a simulated pricing duopoly reported in 2020, applying group comparisons, Pearson and Spearman correlations, and exploratory sensitivity checks. The mean unprofitable deviation share is 96.78 percent. Initial price exhibits little pooled association with this outcome, whereas proportional price reductions show a positive correlation of 0.255 and positive associations within all ten initial price groups. Punishment duration has a pooled correlation of 0.327, which falls to 0.018 after excluding the lowest initial price group. Restricting comparisons to common deviation prices changes association magnitudes while preserving a high average outcome. These findings distinguish persistent resistance to specified deviations from stable evidence explaining its behavioral origin. The contribution qualifies interpretations of repeated competition by demonstrating that outcome persistence can coexist with composition sensitive associations. Evaluations of automated pricing should therefore compare equivalent deviation opportunities and distinguish observed profitability patterns from attributed strategic mechanisms.

Key Words: Algorithmic Collusion; Automated Pricing; Deviation Profitability; Descriptive Analysis; Sensitivity Analysis

1. Introduction:

Automated pricing creates a competition problem that cannot be resolved by observing high prices alone: independently learning algorithms may sustain pricing patterns under which undercutting becomes unprofitable, yet similar outcomes can arise through different learning processes. Distinguishing the outcome from its explanation matters for evaluating pricing systems. An assessment based on price levels, deviation profitability or punishment duration may yield different conclusions about the same algorithmic environment. This analysis examines that distinction through a descriptive reanalysis of published price deviation scenarios, focusing on whether resistance to deviation remains consistent across initial prices, proportional price reductions and subsequent punishment durations. The closest experimental evidence connects algorithmic pricing to the incentives governing repeated competition. Calvano et al. 2020 show that learning algorithms can develop pricing strategies involving temporary punishment followed by a return to cooperation. Their simulations establish that elevated prices can coexist with responses that make specified deviations unprofitable. Klein 2021 demonstrates that the available discrete prices and the timing of interaction influence whether algorithms develop stable prices or pricing cycles. Together, these findings make the pricing environment and the available actions central to interpretation. They also motivate a distinction between the prevalence of unprofitable deviations and the characteristics associated with that prevalence.

Competing explanations further qualify the interpretation of apparent coordination. Hansen et al. 2021 show that independent algorithms can generate prices above competitive benchmarks through correlated experimentation and misspecified demand learning. Abada and Lambin 2023 identify incomplete exploration as another route to seemingly collusive outcomes. These contributions establish that similar aggregate pricing outcomes need not reflect identical strategic processes. Consequently, an observed association between punishment and deviation profitability requires interpretation within the design that generated it. The present analysis examines measured scenario relationships without equating them with an identified learning mechanism or strategic intent. Three observed characteristics organize the analysis. The initial price defines the position from which an agent deviates. A higher starting price may increase the immediate attraction of undercutting while also increasing the future profits placed at risk by a competitive response. Relative price cut measures the proportional size of the departure from that position. Larger reductions can change both the demand captured and the margin earned, preventing an unconditional prediction about their profitability. Punishment duration records the length of the subsequent response, but duration alone does not determine the discounted loss associated with deviation. Calvano et al. 2021 show that punishment phases can vary in intensity over time, illustrating why response length and deterrent strength are conceptually distinct.

The outcome is the published share of specified deviations that are unprofitable. This measure captures resistance to the deviations examined rather than establishing that every feasible deviation is deterred. Its interpretation also depends on aggregation. A scenario summary represents an outcome frequency at a particular combination of initial and deviation prices; it is not an independent simulation run. Moreover, the published grid contains more undercutting opportunities at higher initial prices. Comparisons across initial prices therefore combine variation in starting conditions with variation in the deviations represented. This structure makes scenario coverage relevant to the interpretation of pooled associations. The resulting research gap is bounded and empirical. The source research establishes deviation profitability and punishment patterns within a simulated duopoly. The question addressed here is how the descriptive relationships among those published quantities depend on the composition of the scenario grid. This analysis does not introduce a new experiment. It contributes a secondary examination that separates a persistent aggregate outcome from relationships that change when comparable deviation opportunities are retained. That distinction qualifies how evidence from published simulation tables can be used to assess algorithmic collusion stability.

The dataset contains 95 matched scenarios extracted from Tables A3 and A4 of the online appendix accompanying Calvano et al. 2020. The evidence falls within the publication scope of 2010 to 2024 and comes from simulation results published in 2020, rather than an annual series spanning that interval. Each observation combines an initial price, a strictly lower deviation price, an average punishment duration and an unprofitable deviation share. The analytical approach uses distributional summaries, comparisons within initial price groups, Pearson and Spearman correlations, and exploratory sensitivity checks based on common deviation prices and alternative sample coverage. These procedures characterize the reported scenario grid without treating its cells as independent experimental replications. This study aims to determine how resistance to specified price deviations varies with initial prices, proportional price cuts and punishment duration, and how those relationships change under comparable scenario coverage. The central research question asks whether a high reported frequency of unprofitable deviations is accompanied by consistent descriptive relationships with these three characteristics. The objectives are to examine outcome variation across initial price levels, assess the association with relative price cuts within those levels, and evaluate the relationship with punishment duration across alternative scenario selections.

The completed analysis reveals a distinction between outcome persistence and association sensitivity. The average unprofitable deviation share remains high under alternative weighting and coverage choices. However, initial price has no uniform relationship with the outcome across the complete grid, and the positive pooled association with punishment duration weakens substantially under one exploratory restriction. Relative price cuts retain a positive association within every initial price group. These findings refine the interpretation of deviation resistance by showing that a persistent aggregate pattern can coexist with composition sensitive explanatory relationships. Their practical relevance concerns the evaluation of simulated pricing systems: assessments should distinguish the outcome being monitored, the deviations included and the behavioral explanation attributed to the observed response. Transferability rests on that analytical distinction, while the numerical findings remain specific to the source simulation environment.

2. Theory, Conceptual Framework and Research Questions Development:

Repeated pricing creates an interdependence between immediate profits and future competitive responses. A price reduction may attract demand, but its value also depends on how the rival subsequently adjusts its price. Repeated game theory therefore locates the stability of coordinated pricing in the balance between the immediate benefit of deviation and the discounted losses that follow. Calvano et al. 2020 demonstrate that independently trained pricing algorithms can develop strategies involving temporary punishment and a subsequent return to cooperation. This analysis applies that incentive compatibility perspective to the observed relationship between price deviation characteristics, punishment duration and deviation profitability. Algorithmic collusion stability is operationalized as the reported share of deviations that are unprofitable. A higher share indicates that the specified deviations less frequently improve the deviating agent's discounted profits. The measure captures resistance to the deviations examined rather than establishing that every possible deviation is deterred. The analytical framework connects this observed outcome with three measured characteristics: the price before deviation, the relative price cut and the average punishment duration. These characteristics describe the initial pricing position, the size of the disturbance and the subsequent competitive response.

The price before deviation establishes the starting point from which an agent evaluates undercutting. A higher initial price can increase the potential benefit of capturing additional demand through a price reduction. It can also increase the value of preserving future profits, strengthening the incentive to avoid a disruptive response. These competing mechanisms prevent a general directional prediction based on the initial price alone. Klein 2021 shows that learned pricing patterns depend on the available discrete prices and the structure of sequential interaction. That evidence supports treating price levels as conditions shaping strategic incentives rather than assuming that higher prices necessarily imply greater stability. This analysis examines the reported share of unprofitable deviations across initial price levels while retaining the corresponding deviation prices. Comparisons must account for the fact that higher initial prices permit more undercutting options within the published price grid. Otherwise, differences attributed to initial prices could reflect differences in the deviations represented. Research question 1 asks: How does the reported share of unprofitable deviations vary across prices before deviation within the observed scenarios?

Relative price cut measures the difference between the initial and deviation prices divided by the initial price. It expresses deviation magnitude on a comparable proportional scale. A larger reduction may increase demand captured from the rival, but it also reduces the margin earned on each unit sold. The resulting competitive response can further change the deviation's discounted value. The relationship between deviation magnitude and stability therefore need not be monotonic. The informational role of experimentation reinforces this ambiguity. Hansen et al. 2021 show that independent pricing algorithms can generate prices above competitive benchmarks through correlated experimentation and misspecified demand learning. Their findings establish an alternative route to elevated prices that does not require the same punishment mechanism. For this analysis, the implication is precise: the magnitude of an observed price cut cannot independently establish either competitive behavior or collusive stability. It must be interpreted alongside the reported profitability outcome.

The descriptive comparison consequently examines whether larger proportional cuts correspond to higher, lower or broadly unchanged shares of unprofitable deviations within each initial price level. This preserves the distinction between deviation magnitude and the pricing position from which it occurs. Research question 2 asks: How does the reported share of unprofitable deviations vary with relative price cut within the observed initial price levels?

Punishment duration represents the average number of simulated periods assigned to the punishment phase following a deviation. In the theoretical mechanism, a longer period of reduced profits can offset the immediate benefit of undercutting. Duration alone, however, does not determine the discounted loss because the severity and timing of the response also matter. Calvano et al. 2021 demonstrate that algorithms operating under imperfect monitoring can develop temporary price wars triggered by deviations or adverse demand conditions. Their evidence supports the relevance of dynamic responses while showing that punishment patterns depend on the informational environment. An observed punishment phase also requires careful interpretation. Abada and Lambin 2023 show that apparently collusive outcomes can arise from incomplete exploration rather than advanced strategic coordination. Accordingly, this analysis treats punishment duration as an observed response characteristic

associated with deviation profitability. It does not interpret duration as an independently assigned treatment or infer strategic intent from its presence.

The descriptive comparison pairs average punishment duration with the reported share of unprofitable deviations for each matching price scenario. It examines whether longer punishment phases consistently accompany greater resistance to deviation or whether similar durations coexist with different profitability outcomes. Research question 3 asks: How does the reported share of unprofitable deviations vary across observed punishment durations?

The resulting conceptual framework comprises three observed explanatory characteristics and one observed outcome. Relative price cut is a calculated measure, while the initial price, punishment duration and unprofitable deviation share are published quantities. The relationships represent descriptive associations. They are evaluated across 95 matched price scenarios using the reported cells in Tables A3 and A4 of the source appendix. The research questions allow increasing, decreasing, nonmonotonic and approximately constant patterns without imposing a directional conclusion.

3. Data:

This analysis establishes the provenance and construction of a dataset containing published simulation summaries of price deviations and subsequent punishment. It defines the analytical sample, documents variable measurement and verifies the consistency of the quantities used for descriptive comparisons.

3.1 Data Sources, Population, and Sample Construction:

The dataset draws exclusively on the online appendix accompanying Calvano et al. 2020, provided through the American Economic Association. The source is the appendix dated May 2020. Table A3, panel b, reports the frequency of unprofitable deviations, while Table A4 reports average punishment duration for corresponding price combinations. These table identifiers, row positions and column positions provide the observation locators. Public access does not establish unrestricted reuse, and publisher terms govern the source material. The access date is retained in the accompanying workbook metadata and is distinct from the source publication date. The empirical setting is a simulated pricing duopoly involving interacting learning algorithms. It has no geographical jurisdiction or observed industrial sector. The stated scope of 2010 to 2024 defines the publication eligibility window; the numerical evidence comes from the source published in 2020. Observations are scenario summaries rather than annual measurements. The dataset therefore forms a cross section of published price combinations, with no calendar observation frequency.

The target population comprises price deviation scenarios generated by the source simulation environment. The accessible sampling frame comprises the populated combinations reported in the two appendix tables. Each analytical observation pairs a price before deviation with a strictly lower deviation price. The underlying simulation sessions are not separately represented in the workbook, so the number of independent runs cannot be inferred from the number of extracted cells. Selection retains every published strict undercutting scenario with both outcome and punishment information. The extraction rule excludes combinations where the deviation price equals the initial price because these represent no price reduction. This rule is documented in the workbook construction record. No subsequent eligibility change was introduced during the present verification.

Table 1: Sample construction

Stage	Initial Price Groups Represented	Scenario Observations	Implemented Rule
Populated source frame	10	105	Populated combinations including the diagonal
Diagonal combinations excluded	10	10 removed	Deviation price equals initial price
Eligible undercutting scenarios	10	95	Deviation price strictly below initial price
Matched and verified sample	10	95	Outcome and punishment matched by both price labels
Final descriptive sample	10	95	Complete values for all analytical measures

The retained coverage is triangular: the ten initial price groups contain respectively 5, 6, 7, 8, 9, 10, 11, 12, 13 and 14 eligible deviations. This structure reflects the available price grid. It does not represent a balanced panel or equal sampling across initial price levels.

3.2 Data Integration, Cleaning, and Quality Assurance:

Integration uses the initial price and deviation price jointly as the matching key. The published frequency of unprofitable deviations is paired with the average punishment duration at the same coordinates. Numerical price labels retain their published precision. Source cells remain separate from the analysis sheet, where formulas retrieve values and calculate relative price cuts. This organization preserves the distinction between extracted inputs and derived measures emphasized by Broman and Woo 2018. Verification identified 95 distinct price keys among 95 retained observations. All retained scenarios have corresponding outcome and punishment values, with no unmatched combinations or increase in observations after matching. Repeated initial prices are genuine scenario groupings because each is paired with a different deviation price.

Missingness is zero for the initial price, deviation price, relative price cut, punishment duration and unprofitable deviation share. No deletion for missing values, interpolation, imputation or replacement with zero was applied. Blank combinations outside the populated source grid were not converted into observations. Range checks confirmed that all retained deviation prices are positive and below their corresponding initial prices. Every outcome lies between zero and one, and every punishment duration is positive. Recalculation of relative price cuts reproduced all 95 workbook values. No observation was removed or altered as an outlier. Valid boundary outcomes, including shares equal to one, remain in the sample. The published initial price frequencies are retained as metadata. They are not scenario sample sizes and are not used as analytical weights. Their

rounded total is 0.99, which is preserved without normalization. Prices are model quantities rather than monetary observations, so no currency conversion, inflation adjustment or temporal aggregation was applied.

3.3 Variable Measurement, Construction, and Validation:

Algorithmic collusion stability is represented by the published share of deviations classified as unprofitable. Higher values indicate that the specified deviations less frequently improve discounted profits. This measure describes resistance to the deviations examined; it does not establish resistance to every possible strategy. Abada and Lambin 2023 show that apparently collusive outcomes can also arise through incomplete exploration, reinforcing the need to distinguish an observed profitability pattern from an inferred strategic mechanism. The explanatory characteristics are the price before deviation, the relative price cut and the average punishment duration. The first is an observed model price, the second is calculated from two published prices, and the third is a published response average. They remain separate measures rather than components of a composite index.

Equation 1 defines the implemented relative price cut:

$$c_i = \frac{p_i^b - p_i^d}{p_i^b}$$

Here, c_i denotes the proportional price cut for scenario i, p_i^b denotes its price before deviation, and p_i^d denotes its deviation price. Multiplication by 100 expresses the ratio as a percentage. The calculation uses rounded published price labels rather than unreported underlying price precision.

Table 2: Variable dictionary

Variable or Indicator	Analytical Role	Definition and Construction	Source and Identifier	Unit and Frequency	Transformation and Direction	Missing Value Rule	Coverage and Limitation
Price before deviation	Explanatory characteristic	Published initial price	AEA appendix, Table A3b, row label	Model price units; scenario level	Untransformed; higher values indicate a higher starting price	None missing	95 scenarios; ten distinct levels
Deviation price	Construction input and matching key	Published price imposed during deviation	AEA appendix, Table A3b, column label	Model price units; scenario level	Untransformed; retained below initial price	None missing	95 scenarios; rounded source precision
Relative price cut	Explanatory characteristic	Initial price less deviation price, divided by initial price	Equation 1 using published labels	Ratio; scenario level	Higher values indicate a larger proportional reduction	None missing; all denominators positive	95 scenarios; derived from rounded prices
Punishment duration	Explanatory response characteristic	Published average punishment length	AEA appendix, Table A4, matching cell	Simulated periods; scenario summary	Untransformed; higher values indicate longer punishment	None missing	95 scenarios; an observed response rather than an assigned treatment
Unprofitable deviation share	Outcome	Published frequency of unprofitable deviations	AEA appendix, Table A3b, matching cell	Proportion; scenario summary	Higher values indicate greater resistance to specified deviations	None missing	95 scenarios; underlying cell denominators unavailable
Initial price frequency	Descriptive metadata	Published frequency attached to each initial price row	AEA appendix, Table A3b, frequency column	Proportion; initial price group	Retained without normalization or weighting	None missing	Ten group values repeated across scenarios

Validation addresses numerical consistency and traceability. It does not impose latent variable reliability tests on directly observed quantities. The calculated ratio has positive denominators throughout, and its values agree with the implemented workbook formulas.

Table 3: Descriptive statistics across retained scenario cells

Measure	Valid Count	Mean	Standard Deviation	Median	Minimum	Maximum
Price before deviation	95	1.8295	0.1058	1.8500	1.6200	1.9700
Deviation price	95	1.6099	0.1309	1.5800	1.4300	1.9300
Relative price cut, percent	95	11.8742	6.7603	11.4458	1.6216	27.4112
Punishment duration, simulated periods	95	6.0465	0.5715	6.2200	4.2700	7.1300
Unprofitable deviation share, percent	95	96.7789	2.8627	97.0000	82.0000	100.0000

Statistics are unweighted across published cells. Standard deviations use the sample convention. For percentage measures, the reported standard deviations are expressed in percentage points. These summaries characterize the retained scenario grid rather than the distribution of underlying simulation sessions.

3.4 Analytical Framework, Identification, and Reproducibility:

The analytical framework is descriptive. It examines how the reported unprofitable deviation share varies across initial price levels, relative price cuts and punishment durations. The quantities of interest are observed distributions and scenario differences, rather than causal effects or population parameters estimated from independent runs. Initial prices define the starting conditions, price cuts describe the imposed deviations, and punishment duration summarizes the subsequent response. This ordering preserves the distinction between conditions preceding a deviation and behavior occurring afterward. Klein 2021 demonstrates that algorithmic pricing patterns depend on the structure of interaction and the available price grid. Accordingly, interpretation remains bounded by the source simulation environment and the deviations represented in its published tables. Comparisons retain both price coordinates because the number and magnitude of available deviations differ across initial price levels. The 95 cells are not treated as independent experimental replications. No regression estimator, causal identification procedure, significance test or model specific sample restriction is introduced. Punishment duration is interpreted as a response associated with the outcome, without attributing an independently identified effect to it. Reproducibility rests on the retained published cells, unique scenario identifiers, table coordinates, variable definitions and visible construction formulas. The accompanying workbook separates source inputs from calculated values, permitting reconstruction of Equation 1 and the descriptive summaries. This traceable computational record addresses the reproducibility principle articulated by Sandve et al. 2013: reported quantities should remain connected to the data and operations that produced them.

4. Methods:

4.1 Research Design and Analytical Scope:

This analysis uses a descriptive secondary analysis of published simulation summaries to examine how resistance to specified price deviations varies with initial prices, proportional price reductions and punishment duration. The unit of analysis is a published scenario defined by a unique combination of initial and deviation prices. The outcome is the reported share of deviations that are unprofitable. Comparisons use variation across scenarios within the source simulation environment rather than variation across countries, firms or calendar periods. The analytical quantities of interest are scenario distributions, group means and descriptive associations. Initial price comparisons address differences in starting conditions. Comparisons within initial price groups examine whether proportional deviation magnitude corresponds to the outcome while retaining the same starting price. Punishment comparisons examine whether longer reported responses accompany greater resistance to deviation. These relationships are interpreted as associations because the extracted summaries do not independently vary each characteristic or identify its causal contribution. The source experiment supplies the economic setting, algorithm configuration and underlying simulation procedures documented by Calvano et al. 2020. The present analysis does not rerun that experiment. Its methodological scope is the transparent comparison of quantities reported in the published appendix. This distinction separates the original simulation design from the secondary descriptive calculations undertaken here.

4.2 Measurement Alignment and Descriptive Specification:

The sample construction in Table 1 and measurement definitions in Table 2 determine the analytical inputs. The primary sample contains 95 scenarios across ten initial price groups. Each scenario has complete values for both price coordinates, relative price cut, average punishment duration and the unprofitable deviation share. All primary summaries and pooled associations use these same observations. Relative price cut is calculated using Equation 1. The outcome remains a proportion during computation and is multiplied by 100 for percentage presentation. Punishment duration remains in simulated periods, while prices retain the model units and precision supplied by the source. The analysis preserves these separate measures without constructing a composite index. Equation 2 defines the arithmetic mean used for the complete sample and each specified subset:

$$\bar{z}_S = \frac{1}{n_S} \sum_{i \in S} z_i$$

Here, $\bar{z}_S$ is the mean of measure z over scenario set S, n_S is the number of scenarios in that set, and $z_i$is the value for scenario i. In the primary analysis, each scenario receives equal weight. Group summaries apply Equation 2 separately to scenarios sharing an initial price. Distributional summaries additionally report valid counts, medians, minima, maxima and standard deviations calculated using a denominator equal to the observation countless one. Equation 3 defines the Pearson correlation used to summarize linear association:

$$r_{xy,S} = \frac{\sum_{i \in S}[x_i - \bar{x}_S]\,[y_i - \bar{y}_S]}{\sqrt{\sum_{i \in S}[x_i - \bar{x}_S]^2 \sum_{i \in S}[y_i - \bar{y}_S]^2}}$$

Here, $r_{xy,S}$ is the correlation between explanatory characteristic x and outcome ywithin scenario set S, $x_i$and $y_i$are their scenario values; and $\bar{x}_S$ and $\bar{y}_S$ are their corresponding means. The outcome is the unprofitable deviation share, and the explanatory characteristic is initial price, relative price cut or punishment duration. Spearman correlation is calculated by applying Equation 3 to the ranks of both measures. Tied values receive their average rank. Reporting both coefficients distinguishes linear association from an association in ordering, which is relevant because the published outcome contains repeated rounded values. Relative price cut correlations are also calculated separately within each initial price group.

4.3 Processing, Numerical Verification and Reporting:

Processing begins with the retained published cells and their source coordinates. The initial and deviation prices jointly identify each scenario and link the outcome to its corresponding punishment duration. The published inputs remain separate from the analysis sheet, where cell references retrieve values and formulas calculate relative price cuts. This separation maintains an auditable distinction between source quantities and transformations, consistent with the spreadsheet organization principles described by Broman and Woo 2018. Verification checks the number of observations, uniqueness of the joint price key, completeness of analytical fields and admissible numerical ranges. Every retained scenario must have a positive initial price, a positive deviation price below that initial price, a positive punishment duration and an outcome between zero and one. Relative price cuts are independently recalculated from the two price coordinates and compared with the workbook values.

The completed checks identify 95 unique retained keys, no missing analytical values and no inadmissible values under these rules. No imputation, interpolation, missing value replacement, winsorization or outlier deletion is applied. Boundary outcomes equal to one remain valid observations. Repeated initial prices are retained because they identify groups of distinct deviations rather than duplicate records. Calculations use Python version 3.12.14 with the standard library. Workbook values are read through zipfile and xml.etree. Element Tree. Arithmetic summaries use statistics, while correlation calculations use centered sums and square root operations. Spearman ranks are constructed from the ordered values with average ranks for ties. Calculations are deterministic and require no random seed. Tables 4 to 6 report the distributional summaries, initial price comparisons and descriptive correlations. No significance thresholds are applied. The scenario cells summarize underlying simulation sessions whose dependence and cell denominators are not represented in the workbook. Correlations therefore describe the published grid without confidence intervals or claims of population level significance.

4.4 Exploratory Sensitivity and Replicability:

Sensitivity analyses address the unequal coverage of deviations across initial prices and the influence of particular scenario groups. These checks were exploratory and were introduced after examining the primary descriptive results. They supplement rather than replace the complete sample. The first check restricts observations to the five deviation prices represented at every initial price: 1.43, 1.47, 1.51, 1.54 and 1.58. This produces a common coverage subset of 50 scenarios. Means and correlations are recalculated using the same definitions as the primary analysis. Comparing this subset with the full sample assesses sensitivity to the composition of available deviations; it does not hold proportional cut magnitude constant. The second check assigns equal importance to the ten initial price groups. It first calculates the outcome mean within each group and then takes the arithmetic mean of those ten values. Comparing this quantity with the primary scenario weighted mean identifies the consequence of giving larger groups more influence. Published initial price frequencies remain metadata and are not used as weights.

The third check excludes the five scenarios at the lowest initial price of 1.62 and recalculates the outcome mean and correlations among the remaining 90 observations. This evaluates sensitivity to that group without classifying its observations as erroneous or removing them from the primary results. Interpretation compares changes in direction and magnitude rather than treating agreement across checks as proof of validity. Replication relies on the workbook's retained source cells, scenario identifiers, source locators, measurement dictionary and visible construction formulas. Equations 1 to 3 and the explicit subset definitions specify the calculations needed to reproduce the reported outputs. This connection between inputs, transformations and results follows the reproducibility principle articulated by Sandve et al. 2013. The resulting method answers the research questions through traceable descriptive comparisons while keeping interpretation bounded by the published scenario summaries.

5. Findings:

This analysis examines how initial prices, relative price cuts and punishment duration correspond to the reported share of unprofitable deviations. The findings resolve three research questions through scenario comparisons, descriptive correlations and exploratory sensitivity checks. Interpretation distinguishes resistance to the specified deviations from evidence about the mechanism sustaining that resistance.

5.1 Sample Characteristics and Outcome Distribution:

The analytical sample contains 95 complete price deviation scenarios across ten initial price levels. Each observation combines an initial price, a strictly lower deviation price, an average punishment duration and a reported share of unprofitable deviations. The evidence consists of published simulation summaries from 2020 within the publication scope of 2010 to 2024. Ten combinations involving unchanged prices were excluded. All retained observations have complete analytical values.

Table 4: Descriptive characteristics of the analytical sample

Measure	Observations	Mean	Standard Deviation	Median	Minimum	Maximum
Initial price, model units	95	1.8295	0.1058	1.8500	1.6200	1.9700
Relative price cut, percent	95	11.8742	6.7603	11.4458	1.6216	27.4112
Punishment duration, simulated periods	95	6.0465	0.5715	6.2200	4.2700	7.1300
Unprofitable deviation share, percent	95	96.7789	2.8627	97.0000	82.0000	100.0000

Table 4 indicates that resistance to the specified deviations is widespread but incomplete. The average reported share of unprofitable deviations is 96.78 percent, while the median is 97 percent. Ten scenario cells report 100 percent, whereas the lowest reports 82 percent. These values are frequencies within published scenario summaries; the overall mean is not an estimate of the percentage of all underlying simulation runs exhibiting collusion. The evidence is consistent with the deviation deterrence patterns documented by Calvano et al. 2020, while retaining variation across the particular deviations examined.

5.2 Initial Prices and Resistance to Deviation:

The relationship between initial price and the outcome is nonmonotonic. Table 5 shows that the mean unprofitable deviation share increases from 91.00 percent at an initial price of 1.62 to 98.43 percent at 1.70. It subsequently fluctuates, reaching 98.60 percent at 1.82 before declining to 95.50 percent at 1.97. Higher initial prices therefore do not consistently correspond to greater resistance to deviation.

Table 5: Outcome variation across initial price levels

Initial Price	Scenarios	Mean Unprofitable Share, Percent	Minimum Share, Percent	Maximum Share, Percent	Mean Punishment Duration
1.62	5	91.00	82.00	100.00	4.726
1.66	6	96.33	95.00	100.00	5.070
1.70	7	98.43	98.00	99.00	5.521
1.74	8	96.75	95.00	99.00	5.733

Initial Price	Scenarios	Mean Unprofitable Share, Percent	Minimum Share, Percent	Maximum Share, Percent	Mean Punishment Duration
1.78	9	97.67	96.00	99.00	5.709
1.82	10	98.60	97.00	100.00	6.150
1.85	11	97.73	96.00	99.00	6.434
1.89	12	96.33	94.00	98.00	6.418
1.93	13	97.31	93.00	100.00	6.534
1.97	14	95.50	88.00	100.00	6.447

The pooled Pearson correlation between initial price and the unprofitable deviation share is 0.089, while the Spearman rank correlation is negative 0.042. Both are close to zero, consistent with the absence of a simple ordered relationship across the complete grid. Their interpretation is bounded by unequal scenario coverage: higher initial prices permit more undercutting combinations. An exploratory comparison retains the five deviation prices shared by every initial price group, producing 50 observations. Under this common coverage, the Pearson correlation rises to 0.409 and the Spearman correlation to 0.285. Research question 1 is therefore resolved as a composition sensitive relationship rather than a uniform price gradient. The distinction between the complete grid and common coverage supports interpreting price levels jointly with the available action set, consistent with the importance of discrete pricing opportunities demonstrated by Klein 2021.

5.3 Relative Price Cuts and Punishment Duration:

Relative price cuts display a positive descriptive association with the unprofitable deviation share. Across all 95 scenarios, the Pearson correlation is 0.255 and the Spearman correlation is 0.264. Larger proportional reductions therefore tend to accompany higher reported frequencies of unprofitable deviations, although the relationship is not strong. The association remains positive within every initial price group. Group specific Pearson correlations range from 0.097 to 0.724, and rank correlations range from 0.145 to 0.718. This consistency indicates that the pooled direction does not arise solely from differences between initial price levels. In the 50 scenarios sharing common deviation prices, the Pearson correlation increases to 0.579. Research question 2 is resolved by a positive descriptive pattern whose magnitude depends on the scenario coverage. This pattern concerns the profitability of specified deviations. It does not independently identify how algorithms learned their pricing behavior. Hansen et al. 2021 demonstrate that correlated experimentation can generate elevated prices through a different learning channel, supporting a distinction between the observed outcome and its strategic explanation.

Table 6: Descriptive associations with the unprofitable deviation share

Explanatory Characteristic	Complete Sample Pearson Correlation	Complete Sample Spearman Correlation	Common Coverage Pearson Correlation	Common Coverage Spearman Correlation
Initial price	0.089	Negative 0.042	0.409	0.285
Relative price cut	0.255	0.264	0.579	0.518
Punishment duration	0.327	0.139	0.595	0.378

Punishment duration has a positive pooled Pearson correlation of 0.327 with the outcome, but its Spearman correlation is smaller at 0.139. The difference indicates that the linear association is more pronounced than the ordering of outcomes across punishment lengths. The relationship is also sensitive to sample composition. An exploratory exclusion of the five scenarios at the lowest initial price reduces the punishment duration Pearson correlation to 0.018 among the remaining 90 observations. Conversely, restricting coverage to the five common deviation prices increases it to 0.595. Research question 3 is therefore resolved as a positive pooled association that does not remain similar across alternative scenario selections. Longer punishment cannot be interpreted as a consistently reliable standalone indicator of greater resistance to deviation. This restrained interpretation accords with Abada and Lambin 2023, who distinguish seemingly collusive outcomes from evidence of sophisticated strategic coordination.

5.4 Sensitivity and Synthesis of the Research Questions:

The overall outcome level changes little when initial price groups receive equal weight. Averaging the ten group means produces an unprofitable deviation share of 96.56 percent, compared with 96.78 percent when every scenario receives equal weight. Restricting the sample to common deviation prices produces 97.06 percent. Thus, the high average frequency of unprofitable deviations persists across these alternative summaries, while the relationships with individual characteristics vary more substantially. The sensitivity checks were exploratory and retain the original sample as the primary descriptive basis. They establish that the average outcome and the explanatory associations have different degrees of stability. They do not supply independent experimental replication or causal identification. Numerical magnitudes and changes across sample definitions provide the basis for interpretation, consistent with the emphasis on context and substantive interpretation articulated by Wasserstein and Lazar 2016. The findings answer the research questions in three respects. Initial price has no uniform relationship with resistance to deviation across the complete grid. Larger proportional price cuts show a positive association within all ten initial price groups. Punishment duration has a positive pooled association that weakens substantially under one exploratory sample restriction. The empirical distinction is that a persistently high unprofitable deviation share coexists with composition sensitive relationships between that outcome and the characteristics of individual scenarios.

6. Discussion:

A high frequency of unprofitable deviations does not imply an equally stable explanation of why deviations fail. Tables 4 and 6 reveal this distinction: resistance to the specified deviations remains high, while its associations with initial prices and punishment duration change across scenario selections. The near zero pooled association with initial price contrasts with the consistently positive association between relative price cuts and the outcome within all ten initial price groups. Equation 1 matters because it distinguishes proportional deviation magnitude from the initial price level. The contribution of this reanalysis is

therefore a separation between the persistence of the aggregate outcome and the sensitivity of its descriptive relationships. This refines the interpretation of the punishment patterns documented by Calvano et al. 2020 without treating a secondary analysis of published summaries as an independent experimental discovery. The positive association between relative price cuts and unprofitable deviation shares suggests that deeper undercutting does not consistently offer a more successful departure from established pricing behavior. Table 6 shows that this relationship survives comparisons within initial price groups, reducing the possibility that its direction arises entirely from differences between starting prices. However, the result concerns the discounted profitability of specified deviations rather than an identified transmission mechanism. The theoretical distinction is consequential: larger cuts can change immediate margins, demand allocation and subsequent responses simultaneously. Hansen et al. 2021 demonstrate that learning processes can also produce elevated prices through correlated experimentation. The present evidence therefore sharpens the distinction between deviation resistance as an observed outcome and strategic punishment as one possible explanation. Equation 1 provides a comparable measure of the disturbance, but does not isolate the behavioral channel generating its consequences.

Punishment duration offers a less stable explanatory signal than its positive pooled correlation initially suggests. In Table 6, its linear association exceeds its rank association, while excluding the lowest initial price group reduces the Pearson correlation from 0.327 to 0.018. This contrast indicates that the pooled relationship depends materially on which scenarios enter the comparison. Duration consequently cannot be interpreted as a sufficient measure of deterrent strength. Repeated interaction theory links deviation incentives to discounted profit losses, which depend on the response profile as well as its length. Calvano et al. 2021 describe punishment phases that vary in intensity over their duration, providing a theoretical basis for distinguishing these dimensions. The present reanalysis adds a descriptive qualification: longer punishment and greater deviation resistance need not display a stable relationship across the reported scenario grid. The triangular scenario structure makes coverage part of the substantive interpretation. Table 5 contains progressively more undercutting opportunities at higher initial prices, so pooled comparisons combine variation in starting conditions with variation in the available deviations. Restricting attention to common deviation prices changes the associations in Table 6 even though the average outcome remains high. This reveals that an apparently weak relationship can coexist with a stronger relationship under comparable scenario coverage. Klein 2021 establishes that the available discrete prices shape learned pricing outcomes; the present analysis extends that concern to the interpretation of summaries extracted from a fixed price grid. Analysts evaluating simulated pricing systems should therefore report both the outcome level and its sensitivity to comparable action coverage. Such reporting separates persistent deviation resistance from patterns produced partly by the composition of the evaluated scenarios.

The international relevance lies in this distinction between outcome monitoring and behavioral attribution, rather than in geographical comparisons that the simulation data cannot support. Across the alternative summaries, the high unprofitable deviation share persists, but the relationships used to interpret it vary. Abada and Lambin 2023 show that seemingly collusive outcomes can arise from incomplete exploration, reinforcing the need to avoid assigning strategic intent from aggregate performance alone. For assessments of automated pricing systems, Tables 4 to 6 support evaluating a documented set of deviations, retaining the proportional measure in Equation 1 and checking whether conclusions change under common action coverage. The theoretical implication is that resistance to specified deviations and evidence identifying its cause are distinct empirical claims. This analysis contributes a reproducible illustration of that distinction, narrowing what published scenario summaries establish while making their substantive information more useful for comparisons across algorithmic pricing environments.

7. Conclusion and Implications:

Evaluating algorithmic competition requires distinguishing persistent outcomes from convincing explanations of their persistence. This analysis shows that widespread resistance to specified deviations can coexist with uneven relationships between starting conditions, the scale of disruption and subsequent competitive responses. Its core contribution is a descriptive distinction between outcome persistence and the sensitivity of explanatory associations to scenario coverage. Larger departures display a consistent directional pattern within comparable starting conditions, whereas the apparent importance of response length changes with sample composition. These findings qualify interpretations of repeated competition by showing that observed deterrence and evidence identifying its cause are separate empirical claims. For managers, the implication is to evaluate pricing systems against documented alternatives rather than infer strategic reliability from aggregate performance. Operational assessments should preserve comparable action coverage, distinguish initial conditions from subsequent responses and disclose how scenario selection changes interpretation. For competition authorities, the evidence supports examining deviation profitability alongside the composition of evaluation scenarios before attributing coordination to a particular behavioral mechanism. The broader relevance is methodological: comparable evaluations can improve the transparency of automated pricing assessments across markets. Such discipline provides a firmer basis for decisions affecting firms and consumers while keeping behavioral attribution proportionate to the evidence.

The published scenario summaries provide a basis for extending analysis to individual simulation trajectories. Access to underlying runs would permit uncertainty estimation that accounts for dependence and would clarify how aggregate patterns emerge across learning histories. Experiments varying response intensity and duration separately could distinguish their contributions to deviation profitability. Broader deviation strategies would assess whether resistance extends beyond the specified price reductions. Replication across alternative algorithms, demand structures and market configurations would establish transferability. These extensions would connect the descriptive distinction identified here to direct tests of behavioral mechanisms and their relevance across technological environments.

References:

1. Abada, I., & Lambin, X. 2023. Artificial intelligence: Can seemingly collusive outcomes be avoided? Management Science, volume 69, issue 9, pages 5042 to 5065. https://doi.org/10.1287/mnsc.2022.4623
2. Broman, K. W., & Woo, K. H. 2018. Data organization in spreadsheets. The American Statistician, volume 72, issue 1, pages 2 to 10. https://doi.org/10.1080/00031305.2017.1375989

3. Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. 2020. Artificial intelligence, algorithmic pricing, and collusion. American Economic Review, volume 110, issue 10, pages 3267 to 3297. https://doi.org/10.1257/aer.20190623
4. Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. 2021. Algorithmic collusion with imperfect monitoring. International Journal of Industrial Organization, volume 79, Article 102712. https://doi.org/10.1016/j.ijindorg.2021.102712
5. Hansen, K. T., Misra, K., & Pai, M. M. 2021. Frontiers: Algorithmic collusion: Supra competitive prices via independent algorithms. Marketing Science, volume 40, issue 1, pages 1 to 12. https://doi.org/10.1287/mksc.2020.1276
6. Klein, T. 2021. Autonomous algorithmic collusion: Q learning under sequential pricing. The RAND Journal of Economics, volume 52, issue 3, pages 538 to 558. https://doi.org/10.1111/1756-2171.12383
7. Sandve, G. K., Nekrutenko, A., Taylor, J., & Hovig, E. 2013. Ten simple rules for reproducible computational research. PLOS Computational Biology, volume 9, issue 10, Article e1003285. https://doi.org/10.1371/journal.pcbi.1003285
8. Wasserstein, R. L., & Lazar, N. A. 2016. The ASA statement on p values: Context, process, and purpose. The American Statistician, volume 70, issue 2, pages 129 to 133. https://doi.org/10.1080/00031305.2016.1154108