Following to Lead: Dominant Followership in Sequential Price Competition

Eunsun Kim*

August 31, 2026

Click here for the latest version

Abstract

A dominant platform might be expected to use its market position to charge more than rivals, even for the same products. Amazon does not. Using synchronized high-frequency cross-platform pricing data, I find that Amazon prices at or below the lowest rival price. Its distinctive role appears instead in how it responds to rival price changes, where I document three patterns. First, after a rival raises its price, Amazon's probability of raising its own price more than doubles, whereas rivals show no detectable response when Amazon moves first. Second, when Amazon responds, it tends to match the rival's new price. Third, matched rival price increases remain in place longer than unmatched ones, while matched decreases show no comparable persistence pattern. I call this strategy dominant followership. In addition to documenting this pattern, I provide a stylized model as a potential explanation. A sequential Bertrand model modified to break ties in favor of the dominant firm shows why dominant followership can raise prices above the simultaneous-move Nash benchmark. Because the dominant firm keeps most demand at price parity, it may prefer matching a rival's higher price to undercutting it. The rival is willing to raise price first when that match is predictable; once the price is matched, rival response friction can help preserve it by limiting later undercutting. Using simulation studies, I show that Q-learning algorithms learn dominant-followership behavior endogenously. The learning result separates initiation from incidence: the rival's algorithm learns to initiate higher prices, nearly tripling the market price, while the dominant firm captures 91 percent of the resulting gains. This speaks to the FTC's Project Nessie allegation: individually optimal responses can teach rival algorithms that raising price is profitable, even without explicit coordination. The main insight of this paper is that market power need not operate through price leadership: it can operate through anticipated response. Modeling such sequential responses is important for capturing the essence of modern algorithmic price competition.

Keywords: **Algorithmic Pricing, Artificial Intelligence, Sequential Price Competition, Demand Asymmetry, Reinforcement Learning**

JEL Classification: **L13, L41, C63, C73, L81**

*University of California Los Angeles, eunsun.kim.phd@anderson.ucla.edu.

1 Introduction

A firm with a demand advantage, one that can retain customers even when it does not offer the lowest price, is ordinarily expected to translate that advantage into a price premium (Shaked and Sutton 1982). The online retail data I study run against this expectation: despite its dominant market position, Amazon is usually priced at or below the lowest rival price for the same product. Does this price pattern reflect strong competition, or something else? In this paper, I show that price levels alone cannot answer this question. Firms selling on digital platforms do not choose prices only once or only simultaneously; they observe rivals' prices and respond over time. A dominant firm's market position may therefore matter not only through the price it posts at a point in time, but also through how it responds to rival price changes.

The price-response channel is becoming more important as pricing itself becomes automated. Pricing algorithms allow firms to automate not only the prices they charge, but also how they respond when rivals adjust prices. These algorithmic price-response functions matter more as AI-driven pricing tools make responses more scalable and adaptive.[1] They have also attracted regulatory attention: the U.S. Federal Trade Commission (FTC) has accused Amazon's internal pricing algorithm "Project Nessie" of guiding the entire market toward higher prices through algorithmic price responses (Federal Trade Commission 2024).[2] The regulatory concern extends beyond Amazon and any single enforcement action. Once responses are systematic, they become part of the environment from which rival algorithms learn, especially when the responder is a dominant firm. If rivals learn to anticipate a particular type of response, dominance need not operate through leadership. I ask whether this response channel can raise market prices even while the dominant firm's own price levels appear competitive in that they match rivals' prices.

The paper makes three contributions. First, I document the response pattern I call dominant followership in high-frequency cross-platform pricing data. Second, I explain when this response channel can raise market prices. Third, I show that learning algorithms can generate the pattern

[1] Among 419 B2B pricing executives and decision-makers, 65 to 85 percent reported that their organizations expected to adopt generative and agentic AI in pricing within one to three years; see McKinsey, "B2B pricing: Navigating the next phase of the AI revolution," April 7, 2026.

[2] A related California enforcement action alleges coordinated price increases in which one retailer raised its price and another subsequently matched the higher price. Unlike the mechanism studied here, which operates through independently chosen price responses and does not require coordination, the California filing alleges vendor-mediated agreements among Amazon and competing retailers. See *People v. Amazon.com, Inc.*, Memorandum in Support of Motion for Preliminary Injunction (Cal. Super. Ct. 2026), pp. 2–3.

on their own, using only profit feedback.

The first contribution is empirical. Using a custom scraper that tracks Amazon's first-party prices against every rival platform selling the exact same product, every two hours over eight months, I document how Amazon and its rivals respond to each other's price changes in more than one million synchronized price observations. Amazon's response unfolds in three parts, together showing when rival-initiated higher prices are more likely to persist. First, it reacts far more often than rivals do: after a rival-initiated price increase, Amazon is 12.9 percentage points more likely to raise its own price within 48 hours, 3.5 times the response probability of other major platforms. When Amazon itself moves first, other major platforms show no detectable response. Second, when Amazon does react to a rival increase, it typically matches rather than undercuts the rival's new price: 40.3 percent of its responses land within a narrow parity band. Third, when Amazon matches a rival's higher price, those rival increases are more likely to persist: matched rival price increases are 15.8 percentage points more likely to remain in place after two weeks than unmatched increases, while matched decreases show no such effect. I call this strategy "dominant followership."

The second contribution is theoretical. Understanding dominant followership requires moving away from the symmetric-firm structure common to standard game-theoretic models of simultaneous-move price competition. I introduce two asymmetries. The first is a demand advantage at price parity: one firm retains a larger share of demand. Amazon is a natural example, allegedly accounting for more than 82 percent of gross merchandise value in the online-superstore market, while generally pricing at parity with the lowest rival price. The second is a response asymmetry: one firm can respond to rivals' price changes more systematically and reliably than another. These asymmetries need not favor the same firm. The price-raising channel works when they make the demand-advantaged firm the one expected to respond more reliably. If the follower role shifts to the other firm, the channel can break down.

The dominant followership documented in the data can be rationalized in a simple Bertrand model with one seemingly minor modification: one firm has a demand advantage at price parity. One way to microfound this assumption is a lexicographic decision rule: consumers choose the lower price, but when prices are equal, they favor the dominant platform. This keeps the model close to homogeneous-products Bertrand competition: consumers still choose the lower-priced firm, so the model rules out a direct price-premium channel. Firm identity matters only when prices are tied.

The key implication is that demand advantage changes the value of matching. When a rival raises price, the dominant firm may prefer matching the higher price to undercutting it, monetizing the demand it already holds at parity. The rival, holding only a minority of tied-price demand, has a stronger incentive to undercut instead.

I develop three versions of the model, with successively longer horizons, to build the intuition. The single-market-clearing model shows how the basic response incentive changes with follower identity: after a rival raises price, demand advantage can make matching privately optimal for the dominant firm, while the rival in the same follower role would rather undercut. The two-market-clearing model then asks when the rival is willing to initiate that higher price. Initiation depends on who responds last: it can be profitable when the dominant firm has the final response and matches, but not when the rival has the final response and undercuts. The infinite-horizon model shows that the resulting dominant-followership path can be sustained when firms alternate pricing turns and one firm has a demand advantage at parity. In addition, the same path becomes easier to support when the rival's own ability to react is restricted. Taken together, the models explain why the rival initiates, why the dominant firm matches, and why the match can last.

The third contribution is computational. The theoretical model explains how dominant followership can be sustained when firms understand the incentives they face. In real-world online retail, however, retailers set prices across large assortments where demand conditions and rival pricing patterns differ across products and evolve over time, making this kind of full understanding difficult to assume. Pricing algorithms need not understand this mechanism or begin with the dominant-followership strategy; they learn from realized market feedback. This raises a practical question: can one firm's individually optimal response become the feedback from which another firm's algorithm learns to raise price? I study this question with Q-learning simulations that separate the learning environment from the learner. The simulations show that the pricing behavior underlying the mechanism can emerge through learning, but only in particular environments. Demand advantage and response asymmetry do different jobs: demand advantage strengthens the dominant firm's learned tendency to match, while response asymmetry shifts the rival's upward moves toward higher-price initiation. The two forces reinforce one another: prices rise most when the demand-advantaged firm is also the reliable responder. When the reliable-responder role is instead assigned to the rival, this reinforcement breaks down and the price falls below either single-asymmetry case.

Strikingly, the algorithm that learns to raise prices and the firm that benefits from that learning need not be the same: the rival's algorithm drives the price increase, while the dominant firm captures most of the reward. This is surprising because the theoretical mechanism depends on the dominant firm's predictable match, suggesting that its own algorithm should be the one that has to learn the mechanism. To identify whose learning matters, I let one firm at a time use Q-learning while the other follows a simpler, myopic best-response pricing rule. The decomposition confirms the reversal: when only the dominant firm learns, the market price remains low; when only the rival learns, it reaches 89 percent of the price level achieved when both firms learn. The intuition is a stark asymmetry in what each firm has to figure out. For the dominant firm, matching a higher rival price is a one-shot decision, requiring no experimentation. For the rival, initiating that higher price is a gamble: it only learns whether the bet pays off after repeatedly absorbing the short-run cost of raising price first and waiting to see how the dominant firm responds. Yet once the rival's algorithm learns that initiation can be profitable, the dominant firm captures 91 percent of the resulting profit gains. This has a direct bearing on cases like the Project Nessie allegation: if a rival's own algorithm learns to raise prices, scrutiny focused only on the dominant platform's pricing algorithm may be looking in the wrong place.

The main insight of this paper is that market power need not operate through price leadership. In markets where firms observe and respond to one another's prices, a dominant firm can affect market prices by becoming a predictable follower. This shifts the central object of analysis from the price a firm posts at a point in time to the sequence of price changes: who moves first, who responds, and how demand is allocated once prices are matched. Across the paper, response plays three roles. In the data, it is the empirical object: who responds to whose price changes, and how. In the model, it changes pricing incentives: a firm's anticipated response can make rival initiation profitable. In the learning analysis, it becomes feedback from which another firm's algorithm learns. Simultaneous-move models miss this response channel because they remove the responder role through which demand advantage shapes response incentives and learning. They therefore fail to approximate long-run outcomes when firms differ in both demand advantage and response friction.

The findings in this paper change how algorithmic pricing should be evaluated. Dominant followership need not look anticompetitive in isolation. The dominant firm is not necessarily leading prices upward, communicating with rivals, or committing to a common pricing rule. It may simply be responding to a rival's price as an individually optimal response. The concern is that, once these responses become systematic, they can enter rival algorithms' learning environments and raise market-wide prices without communication or agreement. The resulting welfare concern is that individually optimal responses can shift surplus from consumers toward firms through a channel that standard price-level screens may miss. Regulators and modelers should therefore pay attention not only to price levels, but also to response patterns, timing, and which firms' algorithms learn from those responses. For a firm weighing whether to adopt a learning algorithm, the results carry a specific caution: it may be your own algorithm that learns to raise prices, while a dominant rival is the one who profits from it.

Related Literature. Empirically, the paper shifts attention in the algorithmic-pricing literature from price levels to price response functions. Existing empirical work primarily studies price levels, margins, and whether algorithm adoption changes pricing outcomes in particular markets, including automobiles, gasoline, housing, online marketplaces, and earlier digital-pricing settings (Sudhir 2001; Kannan and Kopalle 2001; Assad, Clark, Ershov, and Xu 2024; Calder-Wang and Kim 2023; Musolff 2024; Zhao and Berman 2025). Less is known about algorithmic responsiveness: whether firms react to rivals' price changes within a given response window, and how those reactions vary across firms with different degrees of demand advantage. A related e-commerce literature documents price dispersion, Buy Box competition, and algorithmic repricing within online marketplaces, especially Amazon Marketplace and online retail (Cavallo 2018; Chen, Mislove, and Wilson 2016; He, Hollenbeck, and Proserpio 2022; Lam 2023; Hanspach, Sapi, and Wieting 2024; Musolff 2024). The Project Nessie allegations point to a different empirical object: cross-platform price responses, in which one online retailer's price change may affect pricing decisions by other retailers. Studying this margin requires observing identical products across platforms at high frequency, which is difficult because Amazon's prices are relatively visible while rival retailers' prices are harder to observe. I address this gap by constructing synchronized high-frequency pricing data that measure how price changes propagate sequentially across retailers.

Theoretically, the paper builds on models of algorithmic pricing and learning. Canonical and recent studies show that independently learning algorithms can generate supracompetitive outcomes in repeated pricing environments with simultaneous or concurrent price choices (Calvano, Calzolari, Denicolò, and Pastorello 2020; Banchio and Mantegazza 2023; Asker, Fershtman, and Pakes 2024; Wang, Huang, Singh, and Srinivasan 2025). These settings are less suited to posted-price markets, where prices are observable and firms can react to rivals' prices at different times, with demand realized between price adjustments. I shift attention from algorithmic collusion in simultaneous or symmetric pricing environments to algorithmic responsiveness under asynchronous price adjustment among asymmetric firms.

The closest related work relaxes symmetry in algorithmic responsiveness. Brown and MacKay (2025) show that speed and commitment on the supply side can allow one algorithmic firm to induce supracompetitive prices from a slower or myopic rival. The dominant-followership mechanism differs because response asymmetry matters through its alignment with demand advantage. Response asymmetry alone is not enough: the price-raising effect depends on whether the reliable responder is also the demand-advantaged firm, and can break down when these roles are misaligned. The anticipated match is credible not because the firm commits to a rule, but because matching is incentive-compatible given its demand advantage at price parity.

This focus on response roles also connects the paper to dynamic oligopoly. Maskin and Tirole (1988) provide a canonical model of alternating-move price competition, in which firms condition pricing decisions on rivals' current prices, and Klein (2021) extends this sequential pricing framework to Q-learning agents. My framework builds on this sequential structure but changes the economic role of moving second: when the follower has a demand advantage, moving second can be more valuable than leading because the follower's predictable match makes rival initiation profitable.

The paper also relates to the empirical literature on price leadership. Byrne and de Roos (2019) show that dominant firms can use leadership and price experiments to create focal points and coordinate market prices. In contrast, the empirical patterns in this paper motivate a focus on how dominant firms respond, rather than how they lead.

Dominant followership also differs from two forms of matching studied in earlier work. Classic repeated-game work studies reciprocal strategies such as tit-for-tat, in which cooperation is sustained because firms punish deviations from past behavior (Axelrod and Hamilton 1981). The price-matching-guarantee literature studies explicit policies under which firms commit to match rivals' lower prices (Edlin 1997; Corts 1997; Hviid and Shaffer 1999; Chen, Narasimhan, and Zhang 2001). Matching in my setting is neither reciprocal punishment nor an advertised commitment. It is an observed algorithmic response pattern whose competitive effect comes from the demand advantage of the firm doing the matching.

Finally, the paper connects algorithmic pricing to the literature on demand-side frictions. Search frictions, product differentiation, loyalty, and switching costs can soften price competition by making consumers less willing or able to switch across firms (Hotelling 1929; Diamond 1971; Varian 1980; Narasimhan 1988; Beggs and Klemperer 1992). This literature mainly studies how such frictions affect price levels and market power. I show that they can also shape the roles firms play in sequential algorithmic competition. In my setting, demand advantage matters not mainly by supporting a price premium, but by changing which responses are profitable and learnable once prices are matched.

2 Empirical Evidence on Dominant Followership

This section builds the empirical case for dominant followership. A platform has a demand advantage when consumers are more likely to choose it over its rivals at matched prices. In a static pricing view, such an advantage creates room to charge a price premium. Amazon is a natural setting to look for this kind of advantage: it is the dominant platform in the online-superstore market defined in the FTC complaint, allegedly accounting for more than 82 percent of gross merchandise value in 2022 (Federal Trade Commission 2024). Yet Amazon is not typically the high-price platform; if anything, it is usually priced at or below its rivals. This is the empirical puzzle: why would the platform best positioned to leverage a demand advantage appear to compete so aggressively on price?

I approach this puzzle by asking where Amazon's market role appears if not in its price level. The key possibility is that higher-price episodes are shaped not only by the price increase itself, but also by the responses that follow. I therefore organize the empirical evidence around three diagnostic questions. First, I ask whether price responses are asymmetric: whether Amazon reacts to rival increases more strongly than rivals react to Amazon or to each other, and whether those responses are concentrated after the rival move (asymmetric responsiveness). Second, I ask where

Amazon's price lands when it responds to a rival increase (matching at parity). Third, I ask whether rival increases last longer when Amazon matches them than when it does not (persistence of matched increases).

These patterns are not what one would expect from frequent repricing, a common shock with delayed Amazon adjustment, or mechanical price chasing. Those accounts would not naturally predict responses concentrated at Amazon and shortly after rival moves, repeated matching around parity, or greater durability precisely when Amazon responds by matching.

I examine these patterns in turn, using the same evidence to assess these alternatives.

2.1 High-Frequency Cross-Platform Pricing Data

Studying price responses requires data that capture a part of competition usually hidden in standard pricing data: the order in which prices change across competing platforms. To recover this order, I build a high-frequency scraping system that tracks identical products across the full set of identified platforms selling each matched variant, collecting prices within synchronized time windows.

Three design features are central. First, prices for a given product are collected across platforms within the same time window, at roughly two-hour frequency. This synchronization makes it possible to compare prices for the same product across platforms at nearly the same moment. In 91.3% of price-change bins, exactly one eligible platform changes price, so the panel can recover an order of moves rather than only comparing prices at a point in time.[3] Second, for Amazon and Walmart, where the scrape observes marketplace seller identity, I restrict observations to first-party platform listings. This is important because the paper studies platform pricing, not seller pricing within a marketplace. For the other major platforms, the scraped pages do not provide a separate marketplace seller field, so I treat the observed listed price as the platform price. Third, I match products at the exact variant level, including SKU, model specifications, and color, so observed price differences reflect platform pricing rather than product heterogeneity.

The panel covers a broad set of consumer electronics and home-appliance categories, including headphones, wearables, smart-home devices, and vacuum and floor-care products, from November

[3] I define a price update as a change from the platform's most recent observed price for the same product. To avoid treating long gaps as immediate responses, I use only cases where the prior observation occurs within the preceding one to six hours.

2024 to June 2025.[4] I begin with prominent Amazon-listed products within these categories and keep only products also sold by at least one non-Amazon platform; because the relevant competitors differ across products, the resulting platform set varies and spans both broad-selection online superstores and narrower specialty or brand-specific retailers. After dropping observations with missing prices or marked out of stock, the resulting platform-pricing panel contains 216 matched products across 52 identified platforms, 792 product-platform pairs, and about 1.02 million in-stock price observations. Products are observed on 3.67 platforms on average, and the average scrape interval is 2.3 hours. The product set covers a broad range of price points and demand levels: median product prices range from below \$50 to above \$500, and Amazon's "Bought in Past Month" demand proxy ranges from 50 to 100,000.

2.2 Empirical Puzzle: Amazon's Price Position

If Amazon's demand advantage translated into pricing power in the usual way, it should rarely be the cheapest option on the page. Figure 1 plots Amazon's price relative to the lowest rival price, using observations in which Amazon and at least one rival list the same product.

Amazon is exactly tied with the lowest rival in 54.2% of comparisons, priced below it in 31.3%, and priced above it in only 14.4%; more than three quarters of comparisons fall within 10% of the lowest rival price. This pattern is stable over time: Amazon's weekly lowest-price share ranges from 77.5% to 88.9% across the sample period, consistently above the weekly mean lowest-price share among major rivals (Walmart, Target, Best Buy, and Home Depot), which ranges from 40.4% to 60.2%.

Amazon's distinctive market role, then, is not expressed through a persistently higher price. The question is whether this reflects strong price competition or whether Amazon's market role appears instead in how it responds when rivals move first.

2.3 Asymmetric Responsiveness to Rival Price Increases

Main response asymmetry. The first empirical pattern is asymmetric responsiveness: Amazon should be unusually likely to raise price after a rival does, and this response should be stronger than the corresponding rival response to Amazon. Table 1 directly compares Amazon's probability

[4] Categories are selected based on sufficient price variation to study adjustment behavior.

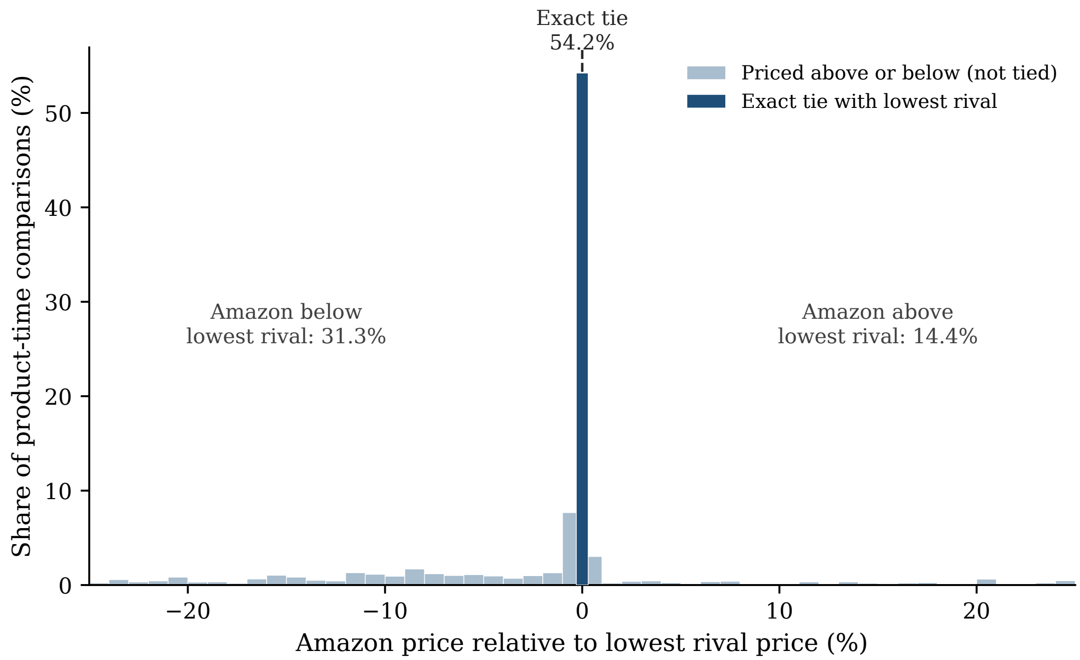

Figure 1. Amazon's Price Relative to the Lowest Rival Price.

Notes: The figure plots Amazon's price relative to the lowest rival price across product-time comparisons in which Amazon and at least one rival have listed prices. A value of zero denotes an exact tie with the lowest rival. Negative values mean Amazon is priced below the lowest rival; positive values mean Amazon is priced above the lowest rival. The figure plots gaps between -25% and 25%; 10.4% of comparisons fall outside this range and are not shown.

of raising price within the next 48 hours, with and without a rival price increase in the preceding period.[5] A prior rival price increase lifts Amazon's probability of raising price within 48 hours by 20.4 percentage points, from 10.2 percent to 30.7 percent. The same pattern remains after controlling for product, platform, week, and day-of-week differences: Amazon is 15.9 percentage points more likely to raise price after a rival increase ($p < .001$; Panel A of Appendix Table A.1). The broad response pattern therefore shows that rival increases predict later Amazon increases.

This broad pattern does not by itself show where the pricing sequence begins. A rival increase may already be a response to an earlier price increase by another platform. I therefore restrict attention to isolated rival increases, in which the ordering is unambiguous: no platform raised price for the product in the prior 48 hours, exactly one identified rival platform then raises price, and

[5] Among the response horizons I consider, the 48-hour window captures a substantial early share of later same-product adjustments in the full price-change panel (33.9%); longer windows show similar lifts but are less tightly tied to the initiating move.

Table 1. Raw Increase Probabilities After Rival Price Increases

Platform	P(Increase within 48h | No rival increase)	P(Increase within 48h | Rival increase)	Difference (p.p.)
Amazon	10.2	30.7	20.4
Walmart	8.5	13.5	4.9
Target	3.5	9.7	6.2
Best Buy	7.5	9.6	2.1
Home Depot	5.9	15.2	9.3

Notes: Entries report each platform's raw probability of raising its own price within 48 hours, with and without a same-product rival price increase at time t.

the responding platform's own price is still unchanged at that moment.[6]

Table 2 establishes this response pattern in three steps. First, Amazon responds strongly to rivals. After an isolated rival increase, Amazon is 12.9 percentage points more likely to raise its own price within the next 48 hours. Second, Amazon responds more strongly than other platforms generally do. The corresponding response among the four major non-Amazon platforms is only 3.7 percentage points, yielding a gap of 9.3 percentage points ($p < .001$). Expanding the comparison to all non-Amazon platforms leaves the conclusion unchanged: their pooled response lift is 4.2 percentage points, and the gap relative to Amazon remains 8.5 points.

Third, rivals do not similarly respond to Amazon. When Amazon is the isolated price increaser, the four major non-Amazon platforms show essentially no response: their response estimate is 0.15 percentage points ($p = .837$). Pooling all non-Amazon platforms gives a similarly small and statistically indistinguishable estimate, effectively zero ($p = .581$). Thus rival increases predict Amazon increases, but Amazon increases do not predict rival increases.

Revenue stakes. A response-channel account of Amazon's responsiveness has a sharper implication: Amazon's response should be stronger when more revenue is at stake. Table 3 tests this implication by splitting products into terciles of Amazon demand times Amazon price. The response estimate is small among low-stakes products, 3.2 percentage points, and much larger among middle- and high-stakes products, 13.4 and 15.7 percentage points. The corresponding estimates

[6] I use the same 48-hour horizon for the quiet window so that the pre-event isolation rule mirrors the post-event response window: shorter quiet windows can treat closely sequenced moves as isolated, while longer quiet windows exclude more events by requiring a longer period with no prior same-direction price change. The analogous definition applies to isolated rival decreases.

Table 2. Responses to Isolated Rival Price Increases

	Price increase
Response to isolated rival price increase	
Amazon	12.94
	(1.79)
Non-Amazon major platforms	3.65
	(0.95)
Difference: Amazon − non-Amazon	9.29
	(1.73), $p < .001$
Observations	597,334
All-platform robustness	
All non-Amazon platforms	4.24
	(0.95)
Difference vs. Amazon	8.54
	(1.77), $p < .001$
Total observations	856,436
Response to isolated Amazon price increase	
Non-Amazon major platforms	0.15
	(0.73), $p = .837$
	$N = 406{,}877$
All non-Amazon platforms	−0.36
	(0.65), $p = .581$
	$N = 665{,}979$

Notes: Entries are percentage-point effects on increasing price within 48 hours after an isolated price increase, relative to otherwise comparable periods without one. The Amazon row uses isolated non-Amazon increases facing Amazon; the non-Amazon row pools Walmart, Target, Best Buy, and Home Depot when they face isolated increases by their own rivals. The difference row reports the estimated difference between these two coefficients and tests whether that difference equals zero, with the standard error clustered by product in parentheses. The top-panel regression contains 597,334 responder-platform observations: 190,457 for Amazon and 406,877 for the four major non-Amazon platforms. The all-platform robustness block expands the non-Amazon comparison to every non-Amazon platform in the panel, giving 856,436 observations: the same 190,457 Amazon observations and 665,979 non-Amazon observations. Observation counts differ because platform coverage differs across matched products. The bottom panel reverses direction, reporting non-Amazon platforms' response to an isolated Amazon price increase. All specifications include product, platform, week, and day-of-week fixed effects; standard errors are clustered by product.

for major non-Amazon platforms are smaller in every tier (0.7, 5.0, and 3.4 percentage points). Thus, higher revenue stakes are associated with stronger responses generally, but the stakes-linked response gradient is much steeper for Amazon, where the revenue stakes of matching are central to the dominant-followership response channel. This revenue-stakes gradient motivates the model below. Proposition 1 formalizes the same payoff logic: matching becomes more attractive when the follower can preserve more revenue at the matched price.

Table 3. Amazon Responds More in Higher-Stakes Products

Response probability by platform	Low stakes (\$6K–\$170K)	Middle stakes (\$179K–\$600K)	High stakes (\$600K–\$5.0M)
Amazon	3.15	13.44	15.69
	(3.49), $p = .366$	(2.93), $p < .001$	(2.68), $p < .001$
Non-Amazon major platforms	0.65	5.01	3.35
	(1.47), $p = .657$	(2.04), $p = .014$	(1.08), $p = .002$
Products	54	66	68
Observations	126,301	224,255	240,179

Notes: Revenue-stakes tiers are terciles of a product-level proxy equal to median Amazon "Bought in Past Month" demand times median Amazon price, computed over the full matched-product panel. The demand proxy converts Amazon's floor-coded display (for example, "3K+ bought in past month") to its numeric lower bound. The ranges shown are rounded dollar values. Entries are percentage-point effects on the probability that the indicated platform(s) increase price within 48 hours of an isolated rival price increase, estimated jointly for Amazon and the pooled non-Amazon major platforms (Walmart, Target, Best Buy, and Home Depot) within each tier. Standard errors, clustered by product, are in parentheses. All specifications include product, week, day-of-week, and platform fixed effects. Product counts are for the isolated-increase regression sample, not all matched products.

Reverse direction. The same response-strength ranking appears for price decreases, where Amazon's response lift after isolated rival decreases is 12.8 percentage points, compared with 3.2 points for the four major non-Amazon platforms (Appendix Table A.2). Unlike the increase case, however, isolated Amazon decreases also elicit detectable rival responses. The response asymmetry is therefore sharpest for upward price movements, while Amazon's unusually strong responsiveness appears in both directions.

Not just delayed adjustment to a common shock. These response patterns are difficult to reconcile with a pure common-shock explanation in which Amazon adjusts with delay. Such an account predicts two forms of uniformity. First, holding the initiating price increase fixed, response rates should be similar across platforms exposed to the same product-level shock. The data instead show a striking cross-platform asymmetry in response rates. In the isolated-increase specification, none of Walmart, Target, Best Buy, or Home Depot approaches Amazon's response individually (Appendix Table A.1, Panel B). Second, if Amazon is only reacting with delay to a shared shock, it should not matter much whether the first mover is a major rival or another platform. Its response rate should therefore be similar after isolated increases by major non-Amazon platforms and by other platforms. The data again show otherwise. Appendix Table A.3 shows that Amazon's upward response lift is 17.5 percentage points after isolated increases by major non-Amazon platforms, compared with 9.9 percentage points after increases by other platforms. This difference is not an artifact of the raw comparison: controlling for product, week, day of week, and the initiator's price-change size, Amazon's response rate is 15.7 percentage points higher after major-platform initiations (clustered standard error 5.8). The response margin is therefore tied not only to the existence of a rival price increase, but also to which rival moves first.

Robustness. Finally, the asymmetric response pattern is not an artifact of competitor-count composition, response-window choice, or product category. Appendix Table A.4 shows that Amazon's response lift is positive in every competitor-count bin and is largest when five or more rival platforms are observed for the same product. Restricting to isolated rival price increases, Amazon's response lift remains significantly larger than the pooled response lift of other major platforms across alternative response windows: the gap is 4.4 percentage points at 12 hours ($p = .001$), 5.3 at 24 hours ($p = .001$), 9.3 at 48 hours ($p < .001$), and 12.0 at 72 hours ($p < .001$), and the pattern also holds across product categories (Appendix Figure A.1).

2.4 Not Before, but Shortly After: Amazon's Response to Rival Increases

The timing evidence asks whether this asymmetric responsiveness is actually tied to the rival's move. If Amazon were already on an upward path, the same events should predict Amazon increases before the rival move. If the estimate only reflected Amazon's frequent repricing, Amazon's increases should be spread throughout the response window rather than concentrated soon after the rival move.

Figure 2 addresses both timing concerns. Panel (a) shows no comparable positive response in

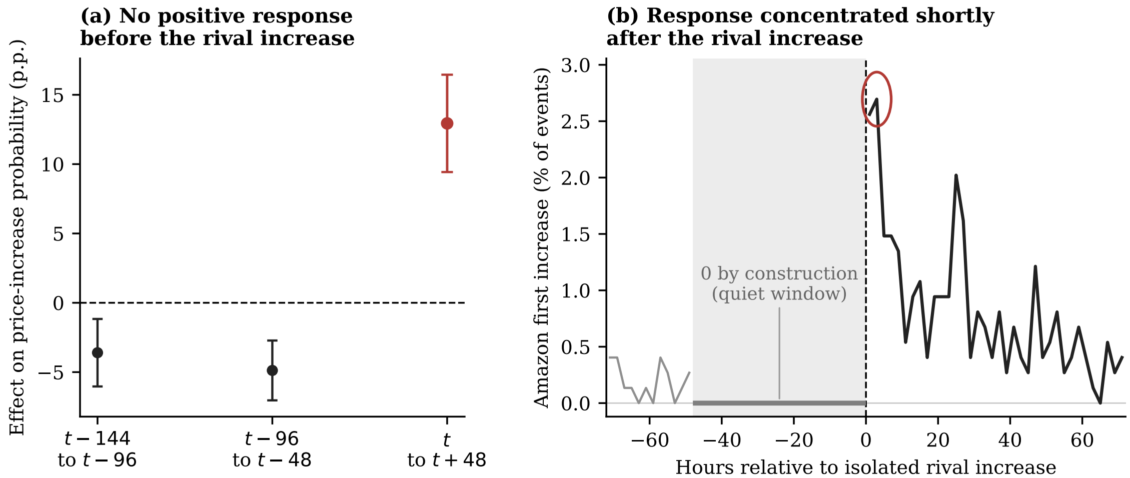

Figure 2. Timing Evidence Around Isolated Rival Price Increases.

Notes: Panel (a) reports Amazon's percentage-point effect on increasing price within 48 hours of an isolated rival increase. Placebo windows assign the same rival-increase events to fake event times 96 or 144 hours earlier. The placebo specifications include product, week, day-of-week, and platform fixed effects, with standard errors clustered by product. Panel (b) conditions on isolated rival increases facing Amazon and plots the share of events for which Amazon's first subsequent price increase occurs in each two-hour bin; the vertical dashed line marks the isolated rival increase. The quiet window is zero by construction because the isolated event requires no same-direction price increase for the product during the prior 48 hours.

the fake earlier windows. Panel (b) plots the timing of Amazon's first price increase around the rival move. Within the main 48-hour response window, 33 percent of Amazon's responses occur in the first 8 hours, with no anticipatory movement beforehand. Amazon is therefore not simply continuing a prior price increase or repricing frequently in general; its increases are concentrated after the rival's observed increase.

2.5 Amazon Matches When Responding to Rival Price Increases

The second empirical pattern is that Amazon often matches the rival's new price. For a rival increase to leave both platforms at the higher price, Amazon's response must reach the rival's new price rather than remain below it. I therefore compare Amazon's response price with that new rival price.

I classify Amazon's first response price within 48 hours as below, at parity with, or above the rival's new price.[7] For example, if Walmart raises the price of a headphone from \$98 to \$109

[7]The parity band is the larger of \$1 or 1% of the rival's new price. Using exact-cent equality instead, exact parity remains the modal response, accounting for 29.4% of cases. The parity-band share is similar under alternative Amazon response windows, holding the 48-hour isolated-event definition fixed: 39.2% at 24 hours, 38.6% at 72 hours,

and Amazon later moves from $98 to $109, I classify Amazon's response as a parity response. If Amazon's response falls outside the parity band, I classify it according to whether it is below or above the rival's new price.

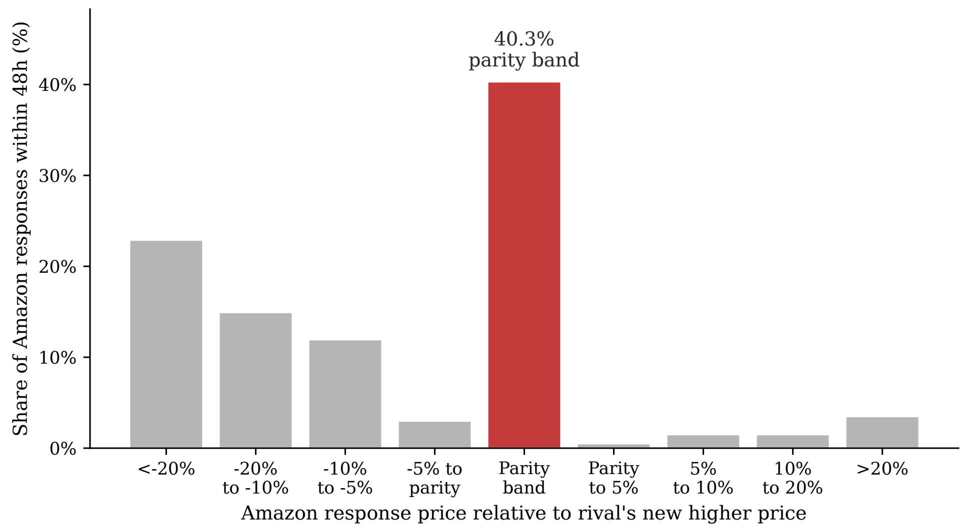

Figure 3. Amazon's Response Price Relative to the Rival's New Price.

Notes: The figure conditions on isolated rival price increases facing Amazon for which Amazon raises price within 48 hours. It plots Amazon's response price relative to the rival's new price. Gaps are measured in percentage points relative to the rival's new price; the parity bar includes responses within the parity band defined in the text.

The response distribution piles up at parity. Amazon's modal response in Figure 3 is the parity band: 40.3% of its responses fall within this band. The matching result is not that Amazon systematically becomes more expensive after rivals raise price. It is that, conditional on responding, Amazon's most common response preserves the rival's higher price by matching it at parity. The shape of the distribution is also informative: a rule designed simply to remain slightly cheaper than the rival would more naturally create a mass just below the rival's new price. Instead, the main mass is at the parity band; responses below parity are often farther below, consistent with Amazon setting a different price rather than slightly undercutting the rival.

The same tendency to land at parity is even stronger after rival price decreases: conditional on Amazon lowering price after an isolated rival decrease, 79.9% of responses fall within the parity band around the rival's new price. The persistence analysis below asks whether this matching

and 38.0% at one week.

behavior has different consequences for increases and decreases.

2.6 Rival Increases Last Longer When Amazon Matches

The third empirical pattern is that Amazon's matching response changes the duration of rival price increases: matched increases last longer. For each isolated rival price change facing Amazon, I compare events in which Amazon reaches the parity band around the rival's new price within 48 hours with events in which it does not. I then track whether the rival's price remains at or above the new higher price after an increase, or at or below the new lower price after a decrease, over the following 60 days. The increase-decrease comparison is useful: if Amazon's matching specifically sustains higher prices, rather than just reflecting general price-following, matched increases should be more likely to stick, while matched decreases should not.

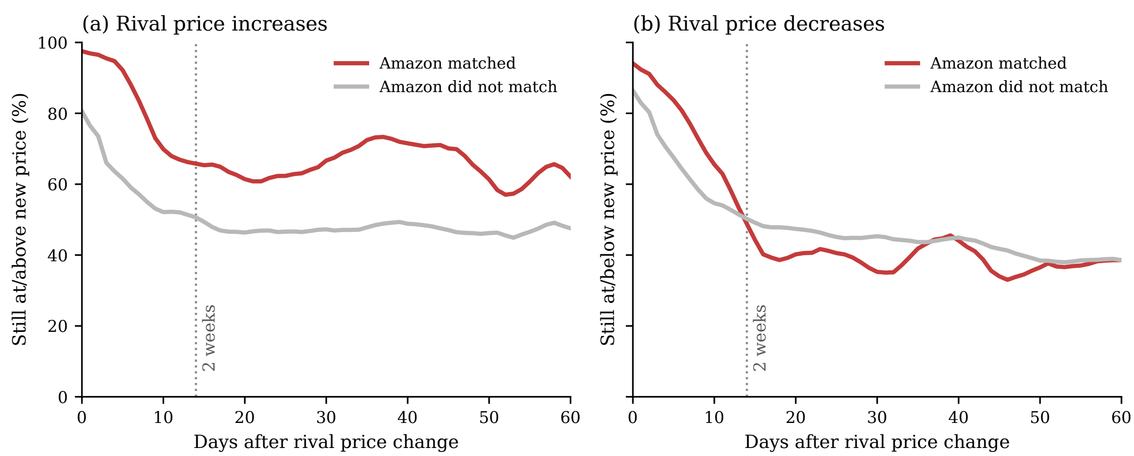

Figure 4. Persistence of Rival Price Changes by Amazon's Response.

Notes: The figure conditions on isolated rival price changes facing Amazon. Panel (a) compares rival increases that Amazon matches within 48 hours with rival increases that Amazon does not match; Panel (b) makes the analogous comparison for rival decreases. Lines report 5-day centered moving averages of daily persistence shares. Persistence means that the rival's later price remains at or beyond the new price: weakly above it after an increase and weakly below it after a decrease. For example, a \$100-to-\$110 increase is persistent only if the later rival price is at least \$110; a \$100-to-\$90 decrease is persistent only if the later rival price is no higher than \$90.

Matched increases persist longer. At two weeks, 67.1% of matched rival increases remain at or above the new price, compared with 51.3% of unmatched increases; product-clustered checkpoint comparisons put the gap at 15.8 percentage points ($SE = 5.8$, $p = .007$). The gap remains similar at 60 days (15.9 percentage points), though it is less precisely estimated ($SE = 8.3$, $p = .058$). Rival decreases do not show the same persistence pattern. If anything, matched decreases persist less often than unmatched decreases at two weeks (44.6% versus 49.2%), though this gap narrows by 60 days (37.7% versus 38.6%). These decrease gaps are not statistically significant (-4.5 percentage points, SE = 5.0, p = .361 at two weeks; -0.9 percentage points, SE = 7.2, p = .895 at 60 days). Amazon's matching response is therefore associated with more durable higher rival prices specifically, not with durable rival price changes in general.

Among active responses, parity matters. The matched-versus-unmatched comparison above pools nonresponses with active responses that fall short of parity. To separate response occurrence from response content, I next condition on cases in which Amazon changes its price within 48 hours. Among isolated rival increases followed by an active Amazon response, increases followed by a parity response are 19.5 percentage points more likely to remain at or above the new price after two weeks than increases followed by an Amazon response price that remains more than 5 percent below the rival's new price (SE = 7.7, p = .013). Thus, the persistence pattern is not driven by any Amazon price change; it is concentrated in responses that move Amazon into the parity band.

Evidence for dominant followership. These results characterize Amazon's dominant-followership strategy through three linked patterns. First, response is asymmetric: Amazon is unusually likely to react after rival price increases, and those reactions are concentrated shortly after the rival move. Second, when Amazon responds, its responses often land at the rival's new price rather than simply undercutting it. Third, matched increases are more durable: when Amazon reaches the rival's new price, the rival's higher price lasts longer. The empirical pattern therefore points to a distinct response channel: a rival raises price, Amazon disproportionately responds, the response often lands at parity, and those parity-matched increases become more durable.

Motivated by this pattern, the next section builds a model to explain why matching responses can be incentive-compatible for a demand-advantaged firm.

3 Why Dominant Followership Can Raise Prices

When two undifferentiated firms compete only on price on a discrete grid, the standard benchmark is the simultaneous-move Bertrand model, which predicts prices near marginal cost.[8] Sequential moves alone do not overturn this intuition: if firms are otherwise symmetric and can undercut each other, the force pushing equilibrium prices toward marginal cost remains. The standard intuition changes once the model is modified in a seemingly minor way: instead of assuming that firms are identical in every respect, I allow one firm, hereafter the dominant firm, to capture more than half of market demand at price parity. This changes the dominant firm's response incentives. The key reversal is that demand advantage can make following more valuable than leading: at parity, the dominant firm keeps enough demand that matching a rival's price increase can beat undercutting it. A rival's higher price can therefore survive not because the dominant firm leads the market upward, but because the dominant firm finds it profitable to follow. I now turn to the model.

3.1 Environment

Suppose there are two firms, indexed by $i \in \{D, S\}$. Firm D is the dominant, demand-advantaged firm, and firm S is the less demand-advantaged rival. Firms choose prices from a finite evenly spaced grid, which reflects the discrete price changes observed in the data.[9]

$$\mathcal{P} = \{0, \delta, 2\delta, \ldots, M\delta\},$$

where $\delta > 0$ is the grid step and $M\delta$ is the maximum price. Let $p_i \in \mathcal{P}$ denote firm i's price. For any x, define the grid floor

$$\lfloor x \rfloor_{\mathcal{P}} = \max\{p \in \mathcal{P} : p \leq x\}.$$

For numerical illustrations, I use the normalized grid $\mathcal{P} = \{0, 0.2, \ldots, 1\}$, so $\delta = 0.2$. In this example, the symmetric simultaneous-move game has Nash equilibria at 0, 0.2, and 0.4.

[8] With marginal cost normalized to zero and prices chosen from the discrete grid, the common-price Nash equilibria of the symmetric simultaneous-move game are 0, δ, and 2δ; the highest is 2δ.

[9] It can also be interpreted as a reduced-form representation of limited consumer price perception: under Weber's law, the smallest noticeable change in a stimulus is proportional to its initial level, so sufficiently small relative price differences may generate the same perceived price and demand response (Fechner 1860; Monroe 1971, 1973). Observed pricing patterns are consistent with discrete, economically meaningful price steps: average dollar adjustment sizes are \$5.17 for products under \$50, \$13.08 for products between \$50 and \$150, \$28.55 for products between \$150 and \$500, and \$81.64 for products above \$500.

Total demand is normalized to one. Demand is allocated according to

$$q_D(p_D, p_S) = \begin{cases} 1, & p_D < p_S, \\ \lambda, & p_D = p_S, \\ 0, & p_D > p_S, \end{cases} \qquad q_S(p_D, p_S) = 1 - q_D(p_D, p_S).$$

The low-price firm receives the entire market. The first primitive is demand advantage, λ: when prices are equal, the dominant firm receives share $\lambda \in [1/2, 1]$, with $\lambda = 1/2$ corresponding to the symmetric benchmark. Relative to homogeneous-products Bertrand competition, the only departure is at price parity: firm identity matters only when prices are tied. The key payoff object is therefore how demand is allocated once prices are matched. This focus is consistent with the data: Amazon is usually priced at or below rivals, and its most important responses are matching responses.

The demand advantage can be microfounded by a simple lexicographic decision rule: consumers first compare prices, but when prices are equal or perceived as equivalent, they choose based on firm identity, brand, membership programs, switching frictions, or other non-price advantages that favor the dominant firm.

This parsimonious demand specification is intentionally conservative: rather than a vertical-differentiation component that would let the dominant platform retain demand above parity, it makes undercutting maximally attractive and removes any direct price-premium channel. Adding such a component would tend to soften competition directly; by omitting it, the model asks whether prices can rise above the symmetric simultaneous-move benchmark even without any direct product-differentiation channel.

Timing and market clearing. The baseline analysis compares two standard timing benchmarks with a single market clearing: simultaneous-move Bertrand competition, in which firms set prices without observing each other's current choice, and sequential-move Bertrand competition (Stackelberg timing), in which one firm, the leader, sets its price first and the other, the follower, observes it before responding. With marginal cost normalized to zero, realized profit is

$$\pi_i(p_D, p_S) = p_i q_i(p_D, p_S).$$

The two-clearing and infinite-horizon extensions below relax the one-shot structure. They introduce a second primitive, response reliability ρ_i, the probability that firm i can revise price after a rival changes price; lower ρ_i corresponds to stronger response friction.[10]

Model intuition. The model isolates three conditions behind dominant followership. First, demand advantage, λ, must be large enough that matching a higher rival price is more profitable for the dominant firm than undercutting it. Second, the dominant firm's response must be predictable enough for the rival to bear the short-run cost of initiating the increase. Third, once the higher price is matched, the rival must face enough response friction that undercutting is no longer attractive. These conditions map to the empirical patterns in Section 2: asymmetric response, matching at parity, and persistence of matched increases. The direction of response asymmetry is therefore central: dominant responsiveness helps create the higher price, while rival response friction helps preserve it.

3.2 Why the Dominant Firm Matches: Single Market-Clearing Model

Matching versus undercutting. The key mechanism can be seen from a firm's best response to a positive rival price $p \in \mathcal{P}$. On the price grid, the relevant local choice is whether to match p or undercut it by one step, $p - \delta$, since any lower price also receives the entire market but yields a lower margin. If the firm receives demand share s at price parity, matching yields sp, while undercutting yields $p - \delta$. The firm is therefore willing to match if and only if

$$sp \geq p - \delta,$$

[10] One possible source of response friction is omnichannel coordination. Retailers that operate both online and physical stores may need to coordinate online prices with store prices, local promotions, inventory, fulfillment systems, customer-service rules, and price-match policies. Fast online repricing may also require capabilities that differ from those used in store-based retailing. I therefore treat ρ_i as a reduced-form measure of a platform's ability to respond to rival price changes, rather than as a claim about a particular internal pricing technology.

or equivalently,

$$p \leq \frac{\delta}{1-s}. \tag{1}$$

The highest rival price that supports matching, $\delta/(1-s)$, is increasing in the firm's demand share at parity.

Simultaneous pricing. In the simultaneous-move game, any pure-strategy equilibrium has $p_D = p_S = p$. Condition (1) requires $p \leq \delta/(1-\lambda)$ for the dominant firm and $p \leq \delta/\lambda$ for the rival. Under demand asymmetry, $\lambda > 1/2$, the rival's condition is more restrictive. Hence the highest equilibrium price is

$$p^{Sim} = \left\lfloor \frac{\delta}{\lambda} \right\rfloor_{\mathcal{P}} < 2\delta,$$

where 2δ is the highest equilibrium price in the symmetric case. Thus, demand asymmetry by itself is not enough to raise prices above the highest symmetric-case equilibrium, and it can even lower the equilibrium price. The equilibrium price is disciplined by the rival, which has the stronger incentive to undercut.

Sequential pricing. I first study a single market-clearing benchmark in which the follower observes the leader's price and can respond with certainty. Sequential pricing changes the constraint because the follower chooses its price after observing the first mover's price. In the single market-clearing game, demand is realized immediately after the follower's response, so the first mover cannot respond again. The highest sustainable matched price is therefore pinned down by the follower's incentive to match rather than undercut, which depends on the asymmetric tie-breaking rule at price parity.

Let p_D^{follow} denote the highest matched price sustainable when the dominant firm follows, and let p_S^{follow} denote the corresponding price when the rival follows.

Proposition 1 (Dominant followership raises equilibrium prices)**.** *Suppose $\lambda \in (1/2, 1)$. The highest matched price sustainable under sequential pricing satisfies*

$$p_D^{follow} > p_S^{follow} = p^{Sim}.$$

Specifically,

$$p_D^{follow} = \left\lfloor \frac{\delta}{1-\lambda} \right\rfloor_{\mathcal{P}} \geq 2\delta,$$

while

$$p_S^{follow} = \left\lfloor \frac{\delta}{\lambda} \right\rfloor_{\mathcal{P}} = p^{Sim} < 2\delta.$$

Table 4 illustrates these role-assignment effects on the price grid for $\lambda = \frac{2}{3}$.

Table 4. Highest Matched Equilibrium Prices on the Price Grid

Demand environment	Simultaneous pricing	Sequential $D \to S$	Sequential $S \to D$
Symmetric demand, $\lambda = 1/2$	2δ	2δ	2δ
Asymmetric demand, $\lambda = 2/3$	δ	δ	$\boxed{3\delta}$

Notes: The price grid is $\{0, \delta, 2\delta, \ldots, 5\delta\}$. Lower matched equilibria may also exist. Because both firms' payoffs are increasing in the matched price, the reported equilibrium is the payoff-dominant matched equilibrium within each timing environment. In sequential pricing, $D \to S$ denotes the dominant firm moving first and the rival following; $S \to D$ denotes the reverse.

Demand advantage changes the value of moving second. Applying the no-undercutting condition (1) to the follower shows why the role assignment matters. A follower with larger tied-price demand is willing to match a higher rival price rather than undercut it. Demand advantage therefore raises prices through the dominant firm's matching incentive, not through a price premium it charges on its own. This logic helps interpret Table 3: Amazon's response to a rival price increase is stronger for higher-revenue-stakes products, where the payoff from a potential match is larger. Appendix Figure A.2 plots these comparative statics for varying values of λ and δ.

Dominant followership, not dominant leadership, raises prices. Table 4 shows that neither demand asymmetry nor sequential timing is enough on its own. With symmetric demand, all timing structures yield 2δ; with asymmetric demand, simultaneous pricing and dominant-firm leadership both leave the price at δ. The price rises only in the role assignment in which the rival moves first and the dominant firm follows. Appendix B.2 formalizes the behavioral separation behind this result and shows that both firms can prefer the dominant-followership role assignment among matched sequential outcomes (Corollaries 1–2). Proposition 5 in Appendix B.8 shows that adding a loyal-consumer component does not eliminate the behavioral separation: the dominant follower can still prefer matching while the rival follower prefers undercutting.

3.3 Why the Rival Initiates a Price Increase: Model with Two Market Clearings

The initiation problem. The single market-clearing benchmark makes initiation costless in a specific sense: if the rival raises price and the dominant firm later matches, demand is realized only after the match, so the rival does not lose current demand while waiting for the response. I next add an intermediate market clearing before the dominant firm can respond, while still allowing demand to be realized again after the response. In this two-clearing game, the rival may lose current demand before being matched, but can recover demand in the second market if the dominant firm later matches. This setup lets me compare two orders that differ only in who holds the final response opportunity: the dominant firm or the rival.

Two-clearing timing. I write the two-clearing game as $i \to j \to M \to i \to M$, where M denotes a market clearing. Firm i first posts a_1, firm j responds with a_2, and the market clears at (a_1, a_2). Firm i then receives a final response opportunity, and the market clears again at (a_3, a_2). Firm k's total payoff is

$$U_k(a_1, a_2, a_3) = \pi_k(a_1, a_2) + \beta\pi_k(a_3, a_2), \qquad k \in \{i, j\}, \tag{2}$$

where $\beta \in (0, 1]$ captures the relative importance of the second market payoff. I solve this finite game by backward induction.

Definition 1 (Price-increase initiation)**.** In the two-clearing game $i \to j \to M \to i \to M$, firm j initiates a price increase if its first response raises price above firm i's inherited price, $a_2 > a_1$.

Costly initiation. The new force is that the rival must bear a short-run cost before the dominant firm can respond. If the dominant firm starts at δ, matching δ gives the rival $(1 - \lambda)\delta$ in both markets, for total payoff $(1 + \beta)(1 - \lambda)\delta$. Initiating a higher price instead means losing current demand, but it can pay off if the dominant firm later matches. If the highest price the dominant firm is willing to match is p^M, initiation gives the rival continuation payoff $\beta(1 - \lambda)p^M$. The rival initiates only when this future matched price is valuable enough to compensate for the first-market loss.

Proposition 2 (Dominant response makes costly rival initiation profitable). *Let p^M denote the highest price that the dominant firm would match if it were given the final response opportunity. From Proposition 1,*

$$p^M = p_D^{follow} = \left\lfloor \frac{\delta}{1-\lambda} \right\rfloor_{\mathcal{P}}.$$

(i) Dominant firm responds last: $D \to S \to M \to D \to M$. *In the normalized-grid region $\lambda \in (2/3, 3/4)$ and $\beta > 1/2$, δ is the dominant firm's equilibrium initial choice. When the dominant firm starts at δ, the rival cannot profitably undercut, so its relevant choice is whether to match δ or initiate the higher price p^M. The rival initiates by raising price to p^M if and only if*

$$\beta(1-\lambda)p^M > (1+\beta)(1-\lambda)\delta, \quad \textit{or equivalently} \quad \beta > \bar{\beta} \equiv \frac{\delta}{p^M - \delta}.$$

The threshold is weakly decreasing in λ because demand advantage raises p^M.

(ii) Rival responds last: $S \to D \to M \to S \to M$. *When the rival has the final response, no matched price above δ is sustained. The rival is willing to match only the lowest positive grid price, δ; at any matched price above δ, it undercuts by one grid step rather than preserving the match.*[11]

The Final Responder Determines Whether Initiation Survives. When the dominant firm responds last, the rival is willing to incur the short-run cost of raising price because it anticipates a valuable match: it raises price above the dominant firm's price, and the dominant firm matches it. When the rival instead holds the final revision opportunity, it later undercuts the very price it initiated. The reason is the response incentive from Proposition 1: with only minority demand at price parity, the rival prefers undercutting to preserving the match. The key reversal is that the higher price is sustained not by sequential timing alone, but by assigning the final response opportunity to the firm whose demand advantage makes matching privately optimal. Table 5 illustrates this contrast.

[11] Appendix B.4 derives equilibrium paths when the rival responds last.

Table 5. Equilibrium Price Path by Final Responder

Final responder	Move 1 (a_1)	Move 2 (a_2)	Move 3 (a_3)
Dominant firm, $D \to S \to M \to D \to M$	δ	3δ	$\boxed{3\delta}$
Rival, $S \to D \to M \to S \to M$	3δ	3δ	2δ

Notes: The displayed paths apply for $\lambda \in (2/3, 3/4)$ and $\beta \in (1/2, 3(1-\lambda))$. In this region, $p^M = 3\delta$. The first row corresponds to Proposition 2, part (i), and the second corresponds to Proposition 2, part (ii), proved in Appendix B.4.

A predictable follower can induce costly initiation. The dominant firm does not need to initiate the price increase itself for its behavior to matter: its predictable response changes the rival's return to raising price first. Proposition 2, part (i), shows that the future matched price must be valuable enough to compensate the rival for the demand it loses when it raises price first. Demand advantage helps by raising p^M, the highest price the dominant firm is willing to match. Starting low is also part of the equilibrium logic: it lets the dominant firm earn current demand while making the rival bear the initiation cost, whereas starting high would instead give the rival an immediate undercutting opportunity. Appendix B.3 provides the payoff comparison.

Figure 5 plots the implied threshold as λ varies. As the dominant firm's demand advantage increases, the highest price it is willing to match rises, making rival initiation easier to sustain.

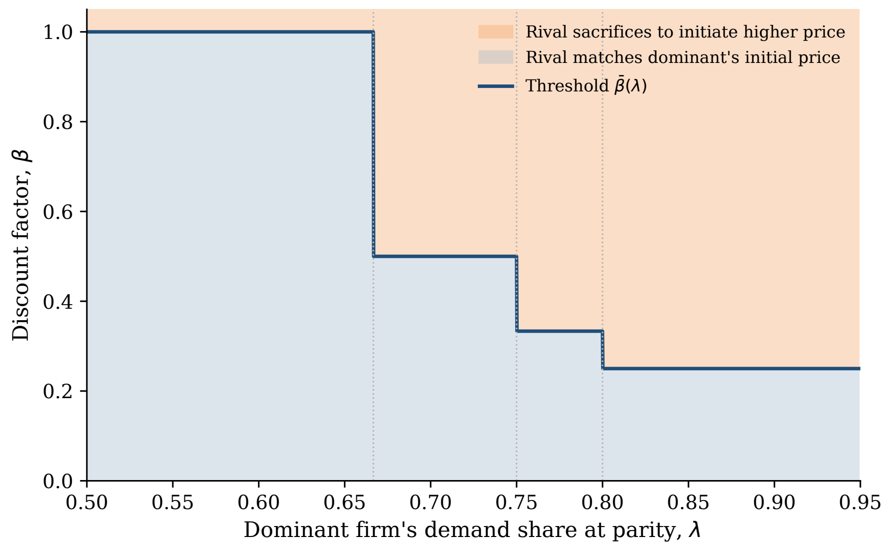

Figure 5. Demand Advantage and Costly Rival Initiation.

Notes: The figure uses the normalized grid $\{0, \delta, 2\delta, \ldots, 5\delta\}$. For each value of λ, it plots the threshold $\bar{\beta} = \delta/(p^M - \delta)$, above which the rival raises price to p^M rather than matching the dominant firm's initial price δ. Dotted vertical lines mark the values of λ at which p^M steps up.

Appendix B.5 allows the dominant firm's final response to arrive sometimes but not always. A higher response probability makes rival initiation more attractive because the rival is more likely to receive the future match.

3.4 Why the Match Persists in an Infinite-Horizon Game

Why Waiting Is Optimal and the Match Persists. The infinite-horizon model asks whether dominant followership can persist when firms keep pricing over time, which requires checking whether the dominant firm prefers waiting to initiating and whether the rival prefers sustaining the high match to undercutting it. The first question matters because the dominant firm benefits from the high-price outcome: if so, why not raise price itself? The second question matters because, once prices are matched, the rival receives only minority demand at parity and may want to undercut.

Repeated-game state. Consider the two-price case $\mathcal{P} = \{L, H\}$, where $0 < L < H$. Firms alternate revision opportunities on a fixed clock, D then S then D then S. The timing is an alternating-move price game in the spirit of Maskin and Tirole (1988): after each scheduled turn, whether or not a price changes, demand is realized at the current posted-price pair. Relative to that benchmark, the model adds demand advantage at parity. The Markov state is (p_D, p_S, m): the current posted-price pair and the firm $m \in \{D, S\}$ whose turn is next.

Candidate policy. The goal is not to classify all Markov-perfect equilibria, but to show that the dominant-followership path that leads to the sustained high-price state (H, H) can be sequentially rational. The candidate policy assigns the costly-initiation role to the rival and the matching role to the dominant firm: from (L, L), the rival raises to H, the dominant firm matches, and both firms remain at (H, H):

price state (p_D, p_S)	$\sigma_D(p_D, p_S)$	$\sigma_S(p_D, p_S)$
(H, H)	H	H
(H, L)	L	H
(L, H)	H	H
(L, L)	L	H

where $\sigma_i(p_D, p_S)$ is firm i's price choice on its pricing turn.

Although the table specifies each firm's action at all four price states, the policy is an equilibrium only if four incentives hold. First, when the rival is low, the dominant firm must prefer setting L to initiating H itself: waiting preserves current demand and leaves the rival to bear the initiation cost. Second, when the dominant firm is low, the rival must prefer setting H to remaining at L: initiating sacrifices current demand but can lead to the future matched price. Third, when the rival is high, the dominant firm must prefer matching at H to undercutting to L, which requires $\lambda H \geq L$. Fourth, when the dominant firm is high, the rival must prefer sustaining or rebuilding H to undercutting to L: the one-period gain from undercutting must not be large enough to justify disrupting the persistent high match.

Formally, these are the four one-period-deviation checks:

$$\text{D waits:} \quad \lambda L + \beta V_D(L, L, S) \geq \beta V_D(H, L, S) \tag{3}$$

$$\text{S initiates:} \quad \beta V_S(L, H, D) \geq (1 - \lambda)L + \beta V_S(L, L, D) \tag{4}$$

$$\text{D matches:} \quad \lambda H + \beta V_D(H, H, S) \geq L + \beta V_D(L, H, S) \tag{5}$$

$$\text{S sustains:} \quad (1 - \lambda)H + \beta V_S(H, H, D) \geq L + \beta V_S(H, L, D) \tag{6}$$

Figure 6 shows the candidate actions and deviations at each posted-price pair.

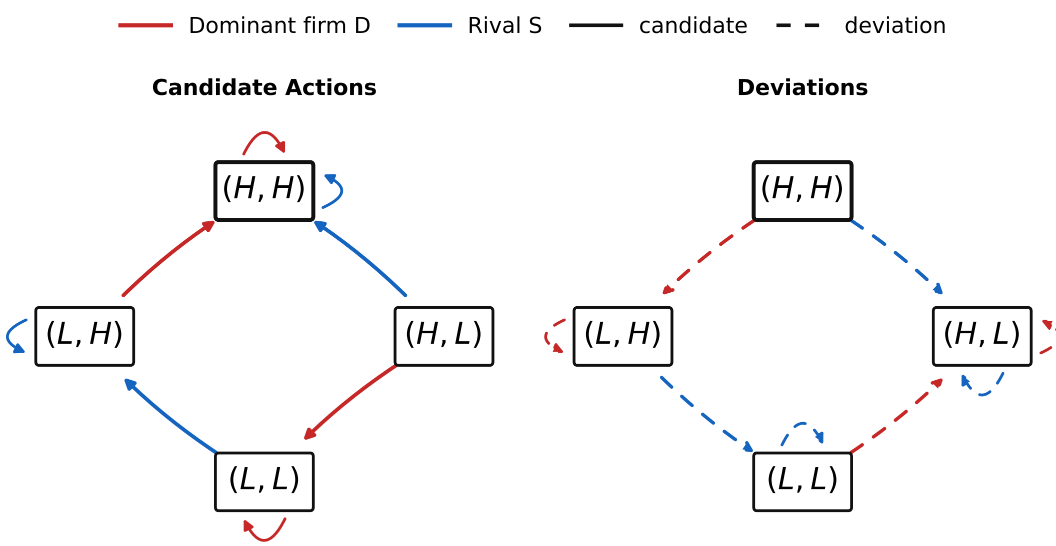

Figure 6. Candidate Dominant-Followership Policy

Notes: States show only the posted-price pair (p_D, p_S). Red arrows denote actions by the dominant firm, and blue arrows denote actions by the rival. Solid arrows are candidate-policy actions; dashed arrows are deviations.

Proposition 3 (Sequential timing can support dominant followership)**.** *Suppose $\lambda > 1/2$, and consider the two-price infinite-horizon game with $\mathcal{P} = \{L, H\}$, $0 < L < H$, alternating pricing turns, and discount factor $\beta \in (0, 1)$. Let $x = L/H$. The candidate policy above is a Markov-perfect equilibrium whenever*

$$\frac{\beta\lambda}{\lambda + \beta} \leq x \leq \min\left\{\beta, \lambda, \frac{(1 - \lambda)(1 + \beta + \beta^2)}{1 + \beta(1 - \lambda)}\right\}.$$

Starting from (L, L), the equilibrium path is $(L, L) \to (L, H) \to (H, H)$, and the market remains at the high matched price thereafter.

Appendix Figure A.4 plots this feasible range for $\beta = 0.9$.

Proposition 3 shows that demand advantage and alternating turns can make deviations costly because returning to the high matched price takes time. The key off-path state is (H, L). If the dominant firm raises first from (L, L), it reaches (H, L) and loses demand while the rival remains low. If the rival undercuts from (H, H), it also reaches (H, L), temporarily destroying the high-price path. In either case, the deviating firm trades a current gain or strategic temptation against the delay required to restore (H, H). This intertemporal cost is what can sustain dominant followership.

Adding response friction. I now add the response-window friction used in the learning simulations (illustrated in Figure 7): after firm j changes price, firm i receives a response opportunity with probability ρ_i. The state becomes (p_D, p_S, m, c), where c indicates whether the preceding period involved a price change and therefore whether a response window is active. I fix $\rho_D = 1$, so the dominant firm faces no response friction, and vary ρ_S, the rival's response probability. This extension asks how rival response friction changes the value of deviations from the same candidate policy.

Proposition 4 (Rival response friction expands the dominant-followership region)**.** *Suppose $\lambda > 1/2$, $\beta \in (0, 1)$, and $\rho_D = 1$. Let $x = L/H$. Let $\underline{x}(\lambda, \beta, \rho_S)$ denote the no-self-initiation threshold defined in Appendix B.7, and let the rival's no-undercutting threshold be*

$$\bar{x}(\lambda, \beta, \rho_S) \equiv \frac{(1-\lambda)\left[1 - \beta^3 \rho_S - \beta^5(1-\rho_S)\right]}{(1-\beta)\left[1 + (1-\lambda)\big(\beta + (1+\beta)\beta^2(1-\rho_S)\big)\right]}.$$

The candidate policy is a Markov-perfect equilibrium, and, starting from (L, L), the market reaches (H, H) and remains there, whenever

$$\underline{x}(\lambda, \beta, \rho_S) \le x \le \min\{\beta,\ \lambda,\ \bar{x}(\lambda, \beta, \rho_S)\}.$$

Relative to pure alternation, lower ρ_S weakly relaxes the no-self-initiation lower bound and raises the no-undercutting upper bound, expanding the parameter region that supports the same dominant-

followership path.

The reason is the same off-path state highlighted above. Lower ρ_S slows recovery from (H, L), so both deviations that pass through that state become less attractive: the dominant firm is less tempted to initiate because the rival may not quickly restore the high price, and the rival is less tempted to undercut because rebuilding the high-price path becomes slower. The proof solves the Bellman equations under this candidate policy and verifies the four incentive checks. Appendix B.7 gives the details.

Appendix Figure A.5 shows how response friction widens this feasible region relative to Figure A.4.

Summary of Theoretical Models. Together, Propositions 1–4 build the mechanism in steps. Proposition 1 shows the basic response incentive: demand advantage raises the value of a future dominant match and makes that match privately optimal for the dominant firm, while the rival in the same follower role would rather undercut. Proposition 2 shows that whether the rival initiates depends on who responds last: initiation can be profitable when the dominant firm has the final response and matches, but not when the rival has the final response and would rather undercut. Proposition 3 shows that, with demand advantage and alternating turns, dominant followership can be sustained in an infinite-horizon game. Proposition 4 then adds response friction and shows that the same dominant-followership path becomes easier to support when the rival's own response probability is lower. Taken together, the models explain why the rival initiates a price increase, why the dominant firm matches, and why the match can last: demand advantage makes dominant matching privately optimal, the identity of the final responder determines whether rival initiation is profitable, repeated interaction can sustain the resulting dominant-followership path, and rival response friction expands the persistence region.

Why Rival Friction Can Help. Taking ρ_i as exogenous does not make response friction economically neutral. The same response friction that constrains the rival can discipline its own incentive to destabilize the outcome. Less response friction need not strengthen the rival's position; it can intensify competition by restoring the rival's ability to undo the high-price outcome.

This interpretation connects back to the response asymmetry in Section 2 and motivates the learning treatment in Section 4, in which the demand-advantaged firm is also the low-friction responder: $\rho_D = 1$ and $\rho_S = 0$.

4 Response as Learning Feedback in Algorithmic Pricing

Section 3 shows how dominant followership can be sustained when firms understand the incentives they face: demand advantage makes matching privately optimal, a sufficiently likely dominant response makes rival initiation worthwhile, and rival response friction helps preserve the matched price. Pricing algorithms, however, need not understand this mechanism or begin with the dominant-followership strategy. They learn from realized market feedback. This creates a different question: when firms must learn how to price, can a response generated by one firm's own profit feedback teach another firm's algorithm that raising price is profitable?

I study this question by separating the learning environment from the learner. First, I hold the learning rule fixed and vary the environment by changing the two forces identified in Section 3: demand advantage and response friction, separately and together. This asks which environments make higher-price initiation and matching emerge in learned pricing behavior. I then hold the dominant-followership environment fixed and vary which firm learns, identifying whose adaptation drives the price increase. Finally, I compare whose algorithm drives higher prices with who captures the resulting gains.

The results reveal a sharp separation between the learning environment, the algorithm that adapts, and the firm that benefits. Learned prices rise the most when the demand-advantaged firm is also the reliable responder. When the rival takes over that reliable-response role instead, the price-raising mechanism breaks down. Yet the reliable-response role is not where the difficult learning occurs. The rival's algorithm must learn to initiate higher prices in anticipation of the dominant firm's response, while the dominant firm captures most of the resulting profit gains. Thus, a response generated by one firm's own profit feedback can shape what another firm's algorithm learns to do, and the firm that learns to raise prices need not be the firm that benefits most from doing so.

4.1 How Responses Enter the Learning Problem

Environment. To translate the theory into a learning problem, I keep the demand environment from Section 3 but remove the assumption that firms understand the underlying incentives. Two firms, D and S, choose prices from a common finite grid $\mathcal{P}$, observe the inherited price pair, and update their pricing policies from realized profit feedback. Each firm's problem is a Markov decision process:

State. $s_{i,t} = (p_{i,t-1}, p_{j,t-1})$, $j \neq i$: the inherited posted-price pair firm i observes when it has a pricing opportunity. Appendix Section A.2 supports this two-price state empirically.

Action. $a_{i,t} \in \mathcal{P}$: the price firm i chooses from the common grid.

Reward. $\pi_i(p_{i,t}, p_{j,t})$: firm i's realized profit under the demand environment in Section 3.1, once both firms' posted prices are set for the period.

Timing. I use the response-window rule introduced in Section 3. Under this timing, demand is realized after each pricing opportunity, reflecting posted-price markets with flow demand between price revisions. A rival price change creates a response opportunity, not an automatic response. If firm i changes price, firm j can revise with probability ρ_j; otherwise, the market clears again at unchanged prices. If firm i does not change price, firm j's ordinary scheduled turn proceeds. Thus, lower ρ_j means stronger response friction. This timing is the learning analogue of the empirical response margin in Section 2: prices adjust sequentially, and allowing ρ_D and ρ_S to differ lets the simulations reflect the asymmetric responsiveness documented there. Figure 7 summarizes the rule.

This sequential design is necessary for studying how responses shape learning. The distinction is not merely whether prices are chosen sequentially: sequential timing creates a feedback link across firms, in which one firm's current action changes the state to which the other responds. The Q-learning target below shows how the resulting response path is credited to the original action. A separate simultaneous-updating benchmark shuts down this response channel entirely: both firms choose prices before demand is realized, so neither observes or responds to the rival's current move. I use it below to isolate the demand effect alone.

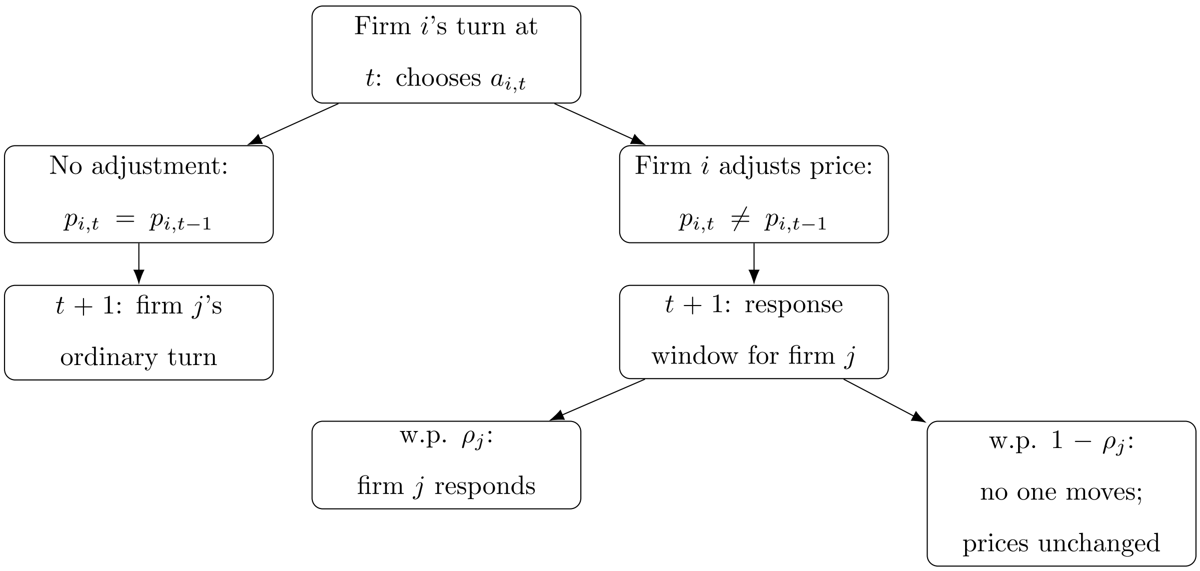

Figure 7. Response-Window Timing.

Notes: The figure illustrates the response-opportunity rule. After firm i's pricing opportunity at t, whether firm j's next scheduled turn is ordinary or gated depends on whether firm i changed its price. In a response window, firm j revises only with probability ρ_j.

Learning rule. Firms use standard Q-learning with ϵ-greedy exploration. Each firm maintains a table $Q_{i,t}(s,a)$, its estimate of the discounted value of price a in state s, and updates it each period as summarized in Algorithm 1.

Algorithm 1 Q-Learning Update for Firm i in Period t

Observe state $s_{i,t} = (p_{i,t-1}, p_{j,t-1})$

Draw $u_t \sim U[0,1]$

if $u_t \leq 1 - \epsilon_t$ **then** ▷ exploit

  $a_{i,t} \leftarrow \arg\max_{a \in \mathcal{P}} Q_{i,t}(s_{i,t}, a)$

else ▷ explore

  $a_{i,t} \leftarrow$ price drawn uniformly at random from $\mathcal{P}$

end if

Post price $p_{i,t} \leftarrow a_{i,t}$

if simultaneous timing **then**

  $\widehat{Q}_{i,t} \leftarrow \pi_i(p_{i,t}, p_{j,t}) + \beta \max_{a' \in \mathcal{P}} Q_{i,t}(s_{i,t+1}, a')$

else ▷ response-window timing

  Accumulate discounted profit over periods $t, \ldots, t+m_i(t)-1$ while $p_{i,t}$ remains posted, as firm j does or does not respond (Figure 7)

  $\widehat{Q}_{i,t} \leftarrow$ accumulated profit + discounted continuation value ▷ two-turn case: Eq. (7)

end if

$Q_{i,t+1}(s_{i,t}, a_{i,t}) \leftarrow (1 - \alpha_t)\, Q_{i,t}(s_{i,t}, a_{i,t}) + \alpha_t\, \widehat{Q}_{i,t}$

The key object is therefore not Q-learning itself, but the feedback attached to a pricing action. In the simultaneous benchmark, a price choice is credited with one period of profit plus continuation value:

$$\widehat{Q}_{i,t}^{sim} = \pi_i(p_{i,t}, p_{j,t}) + \beta \max_{a' \in \mathcal{P}} Q_{i,t}(s_{i,t+1}, a').$$

Under response-window timing, by contrast, a firm's chosen price remains posted while the rival may respond. The profit realized after that response therefore enters the Q-target for the original pricing action. In the usual two-turn case,

$$\widehat{Q}_{i,t}^{rw} = \underbrace{\pi_i(p_{i,t}, p_{j,t})}_{\text{payoff from posted price}} + \underbrace{\beta \pi_i(p_{i,t+1}, p_{j,t+1})}_{\text{profit from rival's response}} + \underbrace{\beta^2 \max_{a' \in \mathcal{P}} Q_{i,t}(s_{i,t+2}, a')}_{\text{continuation value}}. \tag{7}$$

A response by one firm can therefore change the learned value of another firm's earlier price choice, even though neither firm is given the mechanism in advance. Appendix Section C.1 gives the exploration schedule, the exact response-opportunity rule, and the general target for cases in which a response window is not used.

4.2 Simulation Design: Environment, Role Alignment, and Learning

Main simulation treatments. The design separates learning from the environment in which learning occurs. Across the four main treatments, I hold the learning rule fixed and vary only the economic environment: demand advantage and response friction. This isolates how the economic environment shapes learned pricing outcomes while holding the learning technology fixed.

Throughout the 2×2 design, D is the candidate demand-advantaged and reliable-responder firm, while S is its rival. Table 6 reports the resulting 2×2 design. All four cells use the response-window timing described above. The bottom-right cell is distinctive not because it contains more asymmetry, but because the two asymmetries assign compatible strategic roles to the same firm: D has the demand advantage that strengthens the incentive to match in the theory and the reliable response opportunity through which such matching can affect S's learning. This is the dominant-followership treatment.

Table 6. Simulation Treatments

	Response environment	
Demand environment	Symmetric response	Asymmetric response
Symmetric demand	$\lambda = 0.5,\ (\rho_D, \rho_S) = (0.5, 0.5)$	$\lambda = 0.5,\ (\rho_D, \rho_S) = (1, 0)$
Asymmetric demand	$\lambda = 0.7,\ (\rho_D, \rho_S) = (0.5, 0.5)$	$\boxed{\lambda = 0.7,\ (\rho_D, \rho_S) = (1, 0)}$

Notes: In the symmetric-demand row, D and S split demand equally at price parity. In the asymmetric-demand row, D receives share $\lambda = 0.7$ at price parity. In the symmetric-response column, both firms have the same response probability. In the asymmetric-response column, D is the low-friction responder.

A separate simultaneous-updating benchmark shuts down the response channel entirely and compares $\lambda = 0.5$ with $\lambda = 0.7$, holding the price grid and learning rule fixed. This benchmark isolates the demand effect before introducing any response channel.

Role-alignment counterfactual. The 2×2 design establishes whether demand advantage and response friction interact, but not whether their alignment in the same firm is essential. I therefore hold both the degree of demand asymmetry and the degree of response asymmetry fixed and change only the identity of the reliable responder: a role-swap counterfactual that holds demand advantage fixed at D but assigns the reliable-responder role to S instead. This distinguishes a generic response-asymmetry effect from the dominant-followership mechanism: if response asymmetry alone raises prices, responder identity should not matter; if the mechanism operates through the demand-advantaged firm's response incentive, reversing that identity should break it.

Table 7. Role-Alignment Counterfactual

Treatment	Demand advantage	Response environment	Interpretation
Aligned	$\lambda = 0.7$	$(\rho_D, \rho_S) = (1, 0)$	The demand-advantaged firm is the reliable responder.
Misaligned	$\lambda = 0.7$	$(\rho_D, \rho_S) = (0, 1)$	The rival is the reliable responder.

Notes: Both treatments hold demand advantage fixed at firm D, which receives share $\lambda = 0.7$ at price parity. The counterfactual changes only which firm has the high response probability after a rival price change.

One-learner decomposition in the dominant-followership environment. The main treatments identify the environment that generates higher learned prices; the one-learner decomposition asks whose adaptation generates them. Because both firms are learning simultaneously in the two-learner simulations, that comparison alone cannot reveal whose learning is necessary for the resulting price increase. To identify this, I hold the dominant-followership environment fixed, at $\lambda = 0.7$ and $(\rho_D, \rho_S) = (1, 0)$, and let one firm at a time follow a myopic best response instead of Q-learning, maximizing current-period profit against the rival's last posted price without updating continuation values. This does not eliminate price responsiveness; it eliminates learning about how a current price choice changes future response-contingent profits, so the comparison isolates which side of the dominant-followership mechanism requires intertemporal learning. Holding one side fixed also removes the moving-target problem created when two algorithms continually adapt to each other, giving the learning firm a stationary environment with a well-defined solution. Table 8 summarizes the three resulting cases.

Table 8. One-Learner Decomposition: Design (Dominant-Followership Environment)

	Firm S	
Firm D	Q-learning	Myopic best response
Q-learning	$\boxed{\text{Both firms learn}}$	Dominant firm only learns
Myopic best response	Rival only learns	*(excluded: no learning to identify)*

Notes: Each cell fixes one pricing rule per firm. The excluded cell is omitted because, with neither firm learning, there is no learning process left to decompose. The three remaining cells isolate whose learning drives the market price under dominant followership; the boxed cell is the two-learner baseline against which the other two cases are compared.

Outcome measurement. The main outcome is the transaction price, $p_t^{mkt} = \min\{p_{D,t}, p_{S,t}\}$, the price paid by consumers under the baseline demand allocation. I also report firm-level prices and profits to show how learning changes the distribution of gains. Unless otherwise stated, outcomes are averaged over the final 5 percent of periods across simulation runs.

4.3 When Do Learned Prices Rise?

The two forces reinforce one another. The simultaneous-updating benchmark removes the response channel and shows that demand advantage alone actually lowers the learned price: 0.339 under symmetric demand versus 0.274 under asymmetric demand. With response-window timing, Table 9 shows that demand advantage and response asymmetry reinforce each other.

Table 9. LEARNED TRANSACTION PRICES: DEMAND ADVANTAGE AND RESPONSE ENVIRONMENT INTERACT

	Response environment	
	Symmetric response	Asymmetric response
Symmetric demand	0.388	0.533
	(0.043)	(0.021)
Asymmetric demand	0.433	$\boxed{0.660}$
	(0.039)	(0.036)

Notes: The table reports average transaction prices over the final 5 percent of periods in the two-learner Q-learning simulations. The transaction price is $p_t^{mkt} = \min\{p_{D,t}, p_{S,t}\}$. Standard errors, reported in parentheses, are computed across 30 simulation-run means.

Demand asymmetry alone raises the market price only modestly, from 0.388 to 0.433, while asymmetric response timing alone raises it to 0.533. The highest learned price, 0.660, occurs under dominant followership, when the demand-advantaged firm is also the low-friction responder. Why do the two asymmetries reinforce one another? The learned policies reveal that they operate on different sides of the price sequence.

Demand advantage shapes matching; response timing shapes initiation. Panel A shows how demand advantage shapes the response within each response environment. Introducing demand advantage raises D's learned probability of matching S's higher price from 0.152 to 0.242 under symmetric response, and from 0.558 to 0.834 under asymmetric response. This is the learned counterpart of Proposition 1: demand advantage makes preserving a rival's higher price through matching more attractive. Panel B shows how the response environment shapes initiation. When asymmetric response timing is introduced, S's upward moves become much more likely to create a new higher target: the initiation share rises from 0.590 to 0.984 under symmetric demand, and from 0.721 to 0.974 under asymmetric demand. This lines up with Proposition 2: when D occupies the reliable-response role, S's higher-price attempts are more likely to receive response feedback that can make initiation worthwhile.

The two asymmetries therefore do different jobs. Demand advantage makes the dominant firm's

Table 10. Learned Pricing Behavior Across Environments

Panel A. Dominant matching: probability D matches S's higher price

	Demand environment	
	Symmetric demand	Asymmetric demand
Symmetric response	0.152	0.242
Asymmetric response	0.558	0.834

Panel B. Rival initiation: share of S's upward moves that initiate a higher price

	Response environment	
	Symmetric response	Asymmetric response
Symmetric demand	0.590	0.984
Asymmetric demand	0.721	0.974

Notes: The table reports averages over the final 5 percent of periods across 30 simulation runs. Dominant matching is the probability that D exactly matches S's new price after S posts above D's inherited posted price. Rival initiation is the share of S's upward price moves that end strictly above D's pre-move posted price.

learned response more likely to preserve a rival's higher price through matching, while asymmetric response timing shifts the rival's learned upward moves toward initiation. Together, these forces lead the rival to create higher targets more often and the dominant firm to match them more often, moving the learned market price to 0.660.

Under dominant followership, learned play spends more time at high prices. Figure 8 shows where the algorithms spend time after learning, pooling posted-price pairs across simulation runs in the final part of training.

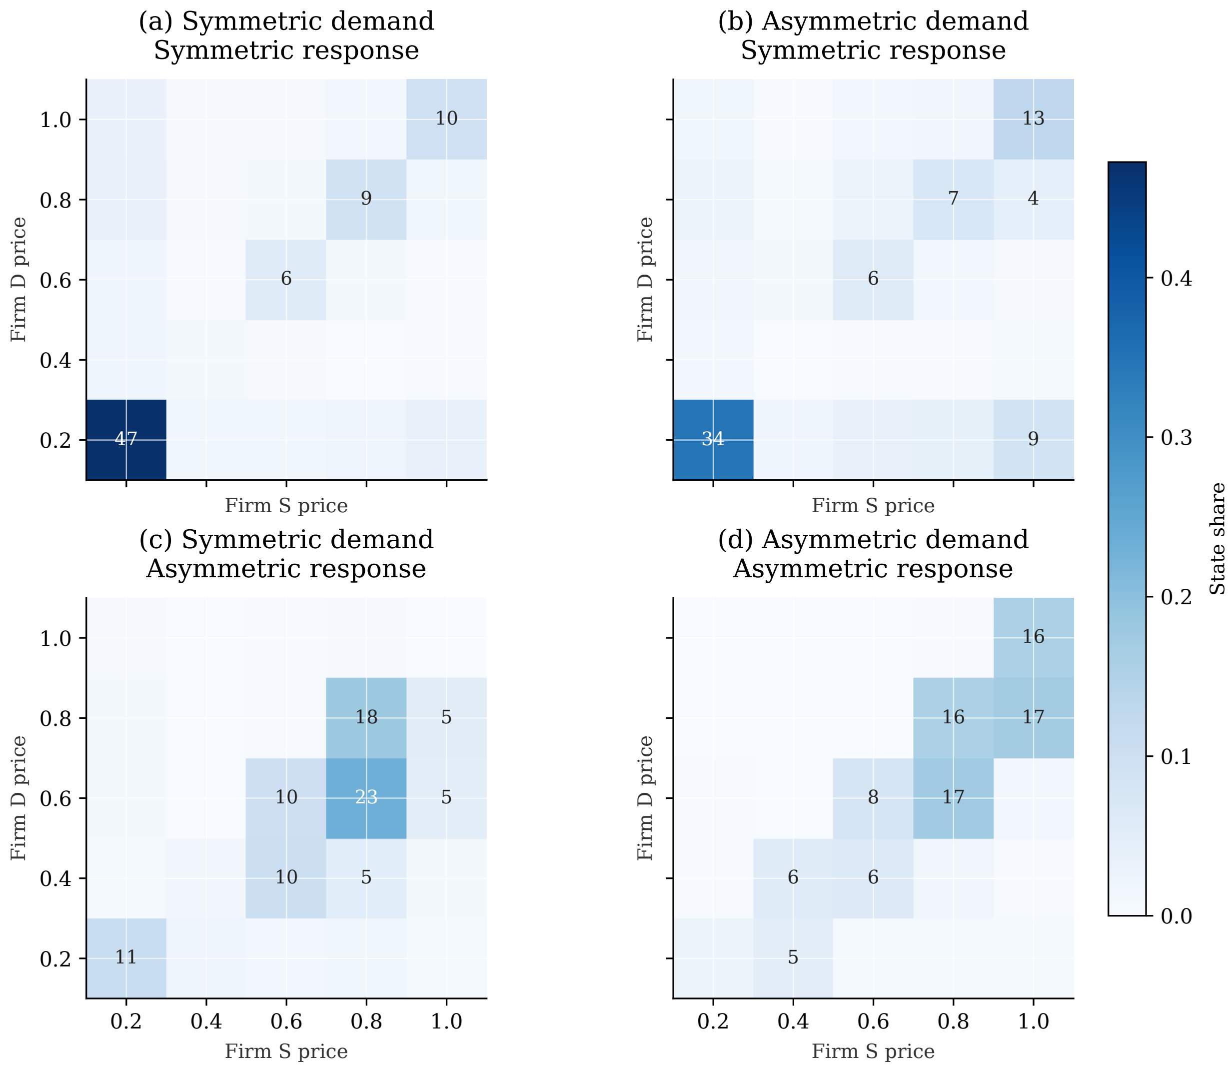

Figure 8. Visited Price States after Learning.

Notes: The figure pools posted-price pairs across all simulation runs in the final 5 percent of training. The vertical axis is firm D's posted price and the horizontal axis is firm S's posted price. Cell labels report percentage shares and are shown for cells with shares of at least 4 percent.

The heatmaps show that the high average prices in Table 9 are not driven by isolated price spikes. The learning process continues to exhibit recurrent price adjustment, but under dominant followership late-period play spends substantially more time in high-price states than in the symmetric-demand, symmetric-response benchmark.

4.4 Misaligned Roles Break the Mechanism

The role-alignment design in Table 7 holds the degree of asymmetry fixed and changes only how the roles are assigned. Demand advantage remains with D, but the reliable-responder role moves from D to S. If prices rise because the two roles are aligned in the demand-advantaged firm, this reassignment should break the mechanism.

Reversing the responder role. I implement this counterfactual by fixing $\lambda = 0.7$, so demand advantage remains with firm D, but reversing the response probabilities: $(\rho_D, \rho_S) = (0, 1)$. In this misaligned treatment, the rival S, not the dominant firm D, is the firm with the high response probability.

Table 11. Aligned versus Misaligned Response Roles

	Aligned $(\rho_D, \rho_S) = (1, 0)$	Misaligned $(\rho_D, \rho_S) = (0, 1)$
Market price	0.660	0.398
Share of periods with a matched high price	0.393	0.021

Notes: Both columns fix $\lambda = 0.7$. The aligned column reproduces the dominant-followership treatment in Table 9. The misaligned column reverses which firm has the high response probability, holding the demand advantage fixed with firm D. Outcomes are averaged over the final 5 percent of periods across 30 simulation runs. A matched high price is a period in which both firms are tied at a price of 0.6 or above.

Misalignment does not merely fall short of the aligned outcome: the learned market price, 0.398, is below either single asymmetry alone (0.433 and 0.533 in Table 9). Response asymmetry raises prices only when its direction supports the dominant-followership role assignment: the demand-advantaged firm must occupy the reliable-response role. When the same response asymmetry is assigned in the opposite direction, the learned price increase is largely eliminated.

How misalignment breaks the learned mechanism. The difference comes from what each firm learns to do. I use each firm's own upward moves to decompose the learned division of labor between initiation and matching.[12] I distinguish three types of upward moves. A *higher-price*

[12] Unlike Table 10, which conditions dominant matching on a prior rival higher-price move, this table conditions on each firm's own upward moves. The two tables therefore answer different questions: Table 10 asks how the

initiation is an upward move that ends above the rival's pre-move posted price, creating a new higher target. An *exact match* ends at the rival's pre-move posted price. A *raise-below-rival* move raises the firm's own price but remains below the rival's pre-move posted price. Table 12 shows that role assignment determines which firm performs each task.

Table 12. Role Assignment and the Learned Mechanism

	Aligned	Misaligned
Upward moves by the rival, S		
Share that initiate a higher price	$\boxed{0.974}$	0.154
Share that exactly match the rival's price	0.025	0.109
Share that raise but remain below the rival	0.001	$\boxed{0.737}$
Upward moves by the dominant firm, D		
Share that initiate a higher price	0.034	$\boxed{0.923}$
Share that exactly match the rival's price	$\boxed{0.861}$	0.033
Share that raise but remain below the rival	0.105	0.044

Notes: Both columns fix $\lambda = 0.7$. A higher-price initiation is an upward move that ends above the rival's pre-move posted price. An exact match ends at the rival's pre-move posted price. A raise-below-rival move raises the firm's own price but remains below the rival's pre-move posted price. Entries are shares of the indicated firm's own upward moves (price increases) in the final 5 percent of periods, averaged across 30 simulation runs.

In the aligned treatment, the learned roles are sharply separated. When the rival S raises price, it almost always creates a higher target: 97.4 percent of S's upward moves are higher-price initiations. When the dominant firm D raises price, it usually matches the rival's posted price exactly: 86.1 percent of D's upward moves are exact matches. The rival tests a higher price, and the demand-advantaged firm repeatedly validates it by matching.

Misalignment does more than lower prices: it reorganizes the learned division of labor. S's upward moves shift away from higher-price initiation, from 97.4 percent to 15.4 percent, and mostly become raise-below-rival moves. D's upward moves shift in the opposite direction, toward higher-price initiation rather than matching. As a result, the market rarely reaches a matched high price:

environment shapes specific learned behaviors, while this table asks what role each firm's upward moves play after responder identity is swapped.

only 2.1 percent of periods are matched high-price states, compared with 39.3 percent in the aligned treatment.

4.5 Whose Learning Drives Higher Prices?

Section 3 shows that the higher-price path depends on the dominant firm's willingness to match rather than undercut. One might therefore expect the dominant firm's algorithm to be the source of the price increase in the learning simulations. The one-learner benchmarks show the opposite. Prices rise when only the rival learns, not when only the dominant firm learns.

The asymmetry comes from how much each firm must discover for this pricing behavior to arise through learning. The dominant firm observes the rival's higher price before choosing whether to match. The rival must raise price first, absorb the immediate market response, and learn whether the dominant firm later makes that higher price profitable. This mirrors Proposition 2's threshold for costly initiation: the rival must learn the long-run return to raising price first, including the continuation value generated by the dominant firm's later response. The one-learner benchmarks isolate each firm's learning problem to identify which side is primarily responsible for the price increase.

Table 13. Learned Prices by Learner Identity (Dominant-Followership Environment)

	Firm S	
Firm D	Q-learning	Myopic best response
Q-learning	$(0.667, 0.776)$	$(0.248, 0.232)$
Myopic best response	$\boxed{(0.592, 0.598)}$	*(excluded)*

Notes: Each cell reports (p_D, p_S), average posted prices over the final 5 percent of periods across 30 simulation runs, under the dominant-followership treatment ($\lambda = 0.7$, $\rho_D = 1$, $\rho_S = 0$). The market price is the period average of $\min\{p_{D,t}, p_{S,t}\}$, which need not equal $\min\{p_D, p_S\}$ computed from the cell averages shown here, since the lower-priced firm can vary across periods. When only one firm uses Q-learning, the other follows a myopic current-period best response; the excluded cell is omitted since, with neither learning, there is no learning process to decompose.

Table 13 shows that learning matters mainly on the initiation side. When only the dominant firm learns, the market price remains low, at 0.225. When only the rival learns, the market price rises to 0.589, close to the two-learner outcome of 0.660. The algorithm that drives the higher price is therefore not the dominant firm's algorithm, but the rival's.

4.6 Who Captures the Gains?

Section 4 identifies whose learning moves prices upward; I now ask who captures the gains from that learning. The rival's algorithm drives the price increase, but it does not capture most of the gains. This is the final asymmetry in the simulation results. Starting from the benchmark in which the dominant firm uses Q-learning and the rival follows a myopic best response, I let the rival switch to Q-learning. Figure 9 shows that this raises the market price by 0.435, from 0.225 to 0.660. Both firms' profits rise, but the gains are highly asymmetric: the dominant firm's profit increases by 0.397, compared with 0.039 for the rival.

Thus, the algorithm that raises prices and the firm that benefits most are not the same. The rival's learning generates higher-price initiation, while the dominant firm's demand advantage allows it to capture most of the resulting surplus. The mechanism therefore has a direct implication for the Project Nessie allegation described in the introduction: a dominant firm need not be the one whose algorithm learns to initiate higher prices in order to capture most of the gains generated by that learning.

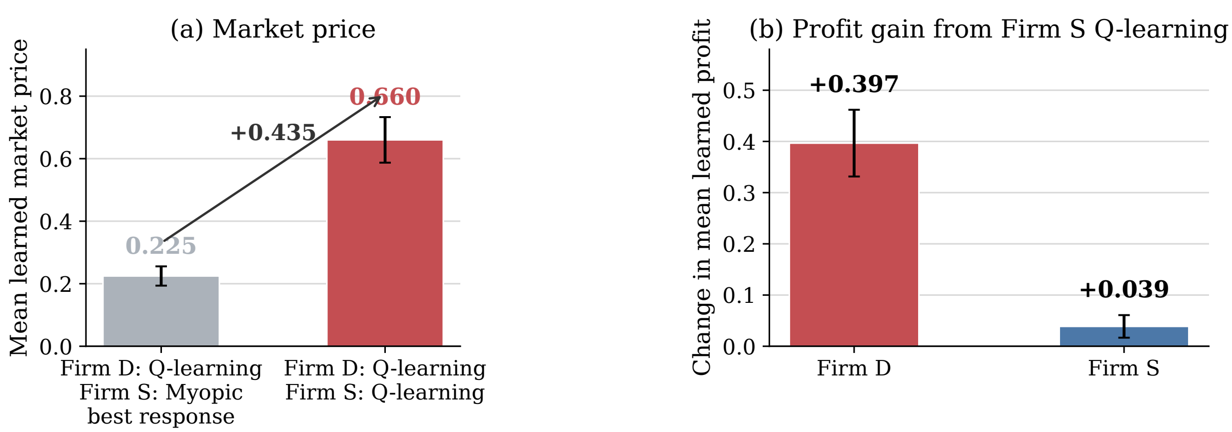

Figure 9. Market-Price and Profit Effects of the Rival's Q-Learning.

Notes: The figure uses the asymmetric-demand, asymmetric-response environment. It compares the benchmark in which the dominant firm uses Q-learning and the rival uses a myopic best response with the two-learner environment in which both firms use Q-learning. Panel (a) reports the change in the market price. Panel (b) reports changes in firm profits. Whiskers report 95 percent confidence intervals across 30 simulation runs. All outcomes are averaged over the final 5 percent of simulation periods.

Appendix C.5 shows that this asymmetric incidence holds at the run level as well: across simulation runs, those in which the rival posts a higher average price than the dominant firm also tend to be the runs in which the dominant firm earns the higher profit premium.

5 Conclusion

Can a dominant firm raise market prices through its responses to rivals' price changes? The central lesson is that it can: market power can operate through anticipated response, not only through leadership. A dominant firm can shape market prices without charging a price premium or initiating the increase itself, if that response changes the profitability of another firm's move.

The paper follows this response channel from behavior to incentives to learning. The data show that Amazon is an unusually systematic responder. The model shows why a demand-advantaged firm may find matching privately optimal, and why that anticipated match can make rival initiation profitable. The learning analysis then shows how the profit realized along a response path can change what another firm's algorithm learns from its own earlier action.

This response channel requires no communication, agreement, commitment, or common pricing rule: each response can be individually optimal given a firm's own demand position. The concern is that once such responses become systematic, they can enter rival algorithms' learning environments and raise market-wide prices without coordination. If price increases can arise through a response system rather than through one firm's algorithm alone, scrutiny confined to a single platform's algorithm may miss the interaction that matters: how one firm's response changes what another firm learns and who captures the gains.

This response channel is not specific to Amazon or to price matching. Its general form is a separation between initiation and validation: one firm takes a risky initial action, and another firm's predictable, privately optimal response changes whether that action pays off. The response could involve product features, fee structures, or other observable strategic choices that rivals learn from over time.

The opening puzzle was that Amazon's apparent price competitiveness sits uneasily with its dominant position. This paper shows why that tension can be misleading. In a sequential pricing environment, a dominant platform can remain at or below rival prices while its anticipated responses affect which rival price increases are profitable to attempt and which higher prices survive.

What looks like aggressive competition in price levels can therefore conceal market power operating through the response the market has learned to expect.

References

Asker, John, Chaim Fershtman, and Ariel Pakes. "The impact of artificial intelligence design on pricing," *Journal of Economics & Management Strategy*, Vol. 33 (2024), pp. 276–304.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market," *Journal of Political Economy*, Vol. 132 (2024), pp. 723–771.

Axelrod, Robert and William D. Hamilton. "The Evolution of Cooperation," *Science*, Vol. 211 (1981), pp. 1390–1396.

Banchio, Martino and Giacomo Mantegazza. "Artificial Intelligence and Spontaneous Collusion," (2023). Working paper.

Beggs, Alan and Paul Klemperer. "Multi-period Competition with Switching Costs," *Econometrica*, Vol. 60 (1992), pp. 651–666.

Brown, Zach Y. and Alexander MacKay. "Algorithmic Coercion with Faster Pricing," Working Paper 34070, National Bureau of Economic Research, 2025.

Byrne, David P. and Nicolas de Roos. "Learning to Coordinate: A Study in Retail Gasoline," *American Economic Review*, Vol. 109 (2019), pp. 591–619.

Calder-Wang, Sophie and Gi Heung Kim. "Coordinated vs Efficient Prices: The Impact of Algorithmic Pricing on Multifamily Rental Markets," (2023). Working paper.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. "Artificial Intelligence, Algorithmic Pricing, and Collusion," *American Economic Review*, Vol. 110 (2020), pp. 3267–97.

Cavallo, Alberto. "More Amazon Effects: Online Competition and Pricing Behaviors," Working Paper 25138, National Bureau of Economic Research, 2018.

Chen, Le, Alan Mislove, and Christo Wilson. "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace," In *Proceedings of the 25th International Conference on World Wide Web* pp. 1339–1349 (2016).

Chen, Yuxin, Chakravarthi Narasimhan, and Z. John Zhang. "Consumer Heterogeneity and Competitive Price-Matching Guarantees," *Marketing Science*, Vol. 20 (2001), pp. 300–314.

Corts, Kenneth S. "On the Competitive Effects of Price-Matching Policies," *International Journal of Industrial Organization*, Vol. 15 (1997), pp. 283–299.

Diamond, Peter. "A Model of Price Adjustment," *Journal of Economic Theory*, Vol. 3 (1971), pp. 156–168.

Edlin, Aaron S. "Do Guaranteed-Low-Price Policies Guarantee High Prices, and Can Antitrust Rise to the Challenge?," *Harvard Law Review*, Vol. 111 (1997), pp. 528–575.

Fechner, Gustav Theodor. *Elemente der Psychophysik*. Breitkopf und Härtel, Leipzig (1860).

Federal Trade Commission. "United States District Court Western District of Washington at Seattle," Case No. 2:23-cv-01495-JHC, Federal Trade Commission, 2024.

Hanspach, Philip, Geza Sapi, and Marcel Wieting. "Algorithms in the Marketplace: An Empirical Analysis of Automated Pricing in E-Commerce," *Information Economics and Policy*, Vol. 69 (2024), p. 101115.

He, Sherry, Brett Hollenbeck, and Davide Proserpio. "The Market for Fake Reviews," *Marketing Science*, Vol. 41 (2022), pp. 896–921.

Hotelling, Harold. "Stability in Competition," *The Economic Journal*, Vol. 39 (1929), pp. 41–57.

Hviid, Morten and Greg Shaffer. "Hassle Costs: The Achilles' Heel of Price-Matching Guarantees," *Journal of Economics & Management Strategy*, Vol. 8 (1999), pp. 489–521.

Kannan, P. K. and Praveen K. Kopalle. "Dynamic Pricing on the Internet: Importance and Implications for Consumer Behavior," *International Journal of Electronic Commerce*, Vol. 5 (2001), pp. 63–83.

Klein, Timo. "Autonomous Algorithmic Collusion: Q-Learning Under Sequential Pricing," *RAND Journal of Economics*, Vol. 52 (2021), pp. 538–558.

Lam, H. Tai. "Platform Search Design and Market Power," (2023). Working paper.

Maskin, Eric and Jean Tirole. “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica*, Vol. 56 (1988), pp. 571–599.

Monroe, Kent B. ““Psychophysics of Prices”: A Reappraisal,” *Journal of Marketing Research*, Vol. 8 (1971), pp. 248–250.

Monroe, Kent B. “Buyers’ Subjective Perceptions of Price,” *Journal of Marketing Research*, Vol. 10 (1973), pp. 70–80.

Musolff, Leon. “Algorithmic Pricing, Price Wars and Tacit Collusion: Evidence from E-Commerce,” (2024). Working paper, Wharton School.

Narasimhan, Chakravarthi. “Competitive Promotional Strategies,” *Journal of Business*, Vol. 61 (1988), pp. 427–449.

Shaked, Avner and John Sutton. “Relaxing Price Competition Through Product Differentiation,” *The Review of Economic Studies*, Vol. 49 (1982), pp. 3–13.

Sudhir, K. “Competitive Pricing Behavior in the Auto Market: A Structural Analysis,” *Marketing Science*, Vol. 20 (2001), pp. 42–60.

Varian, Hal. “A Model of Sales,” *American Economic Review*, Vol. 70 (1980), pp. 651–659.

Wang, Qiaochu, Yan Huang, Param Vir Singh, and Kannan Srinivasan. “Algorithms, Artificial Intelligence and Simple Rule-Based Pricing,” (2025). Working paper.

Zhao, Hangcheng and Ron Berman. “Algorithmic Collusion of Pricing and Advertising on E-Commerce Platforms,” (2025). Working paper.

A Empirical Appendix

A.1 Response Robustness and Heterogeneity

Table A.1. Price Adjustment Responses to Rival Price Changes

	Panel A: Any rival move		Panel B: Isolated rival move	
	Price increase	Price decrease	Price increase	Price decrease
Amazon	15.90	8.91	12.94	12.78
	(3.19), $p < .001$	(3.30), $p = .007$	(1.79), $p < .001$	(1.71), $p < .001$
Walmart	0.82	1.93	4.04	1.30
	(2.11), $p = .697$	(1.96), $p = .324$	(1.86), $p = .030$	(2.02), $p = .521$
Target	3.41	2.37	4.42	4.07
	(1.42), $p = .017$	(1.66), $p = .154$	(1.34), $p = .001$	(1.22), $p = .001$
Best Buy	-0.83	-0.94	1.93	1.47
	(1.29), $p = .521$	(1.53), $p = .539$	(1.34), $p = .149$	(1.21), $p = .224$
Home Depot	4.25	4.37	6.12	8.30
	(5.18), $p = .412$	(4.40), $p = .321$	(2.56), $p = .017$	(3.22), $p = .010$
Product FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
Day-of-week FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Clustered SE	Product	Product	Product	Product
Observations	597,309		597,334	

Notes: The unit of observation is platform-product-time. The responding platform is required not to change price at t. Panel A uses any rival price move for the same product at t. Panel B restricts to isolated rival price moves, in which no platform raised price for the product in the prior 48 hours, exactly one identified rival platform moves at t, and the responding platform's own price is unchanged at that moment. Coefficients are percentage-point effects. Standard errors, clustered by product, are reported in parentheses, followed by the exact two-sided p-value.

Table A.2. Responses to Isolated Rival Price Decreases

	Price decrease
Response to isolated rival price decrease	
Amazon	12.78
	(1.71)
Non-Amazon major platforms	3.17
	(0.92)
Difference: Amazon − non-Amazon	9.61
	(1.59), $p < .001$
Observations	597,334
All-platform robustness	
Amazon	12.69
	(1.72)
All non-Amazon platforms	3.88
	(0.77)
Difference vs. Amazon	8.81
	(1.65), $p < .001$
Total observations	856,436
Response to isolated Amazon price decrease	
Non-Amazon major platforms	2.84
	(0.95), $p = .003$
	$N = 406{,}877$
All non-Amazon platforms	3.45
	(0.80), $p < .001$
	$N = 665{,}979$

Notes: Entries are percentage-point effects on decreasing price within 48 hours after an isolated price decrease, relative to otherwise comparable periods without one. The table mirrors Table 2 for downward price movements. The Amazon row uses isolated non-Amazon decreases facing Amazon; the non-Amazon row pools Walmart, Target, Best Buy, and Home Depot when they face isolated decreases by their own rivals. The all-platform robustness block repeats this comparison using every non-Amazon platform in the pricing panel. The bottom panel reverses direction, reporting non-Amazon platforms' response to an isolated Amazon price decrease. Standard errors are reported in parentheses; for difference rows, the parenthesized value is the standard error of the difference. All specifications include product, platform, week, and day-of-week fixed effects; standard errors are clustered by product.

Table A.3. Amazon Responses by Initiating Rival Type

Initiating rival	Events	Products	Response rate	Response lift
Major non-Amazon platforms	511	133	28.7%	17.5 pp
Other non-Amazon platforms	236	70	20.2%	9.9 pp
Major − other initiators, controlled			15.7 pp	(5.8)

Notes: The sample is Amazon observations after isolated non-Amazon upward first moves. Response rate is the probability that Amazon raises price within 48 hours. Response lift subtracts Amazon's baseline upward adjustment rate for products in the same initiating-rival group when that group did not initiate a clean upward first move. The controlled row estimates the major-versus-other difference within the treated sample, including product, week, day-of-week fixed effects and the initiator's percentage price increase; standard errors are clustered by product.

Table A.4. Amazon Response Lift by Number of Observed Rival Platforms

Competitor count	Events	Response rate	Baseline rate	Response lift
1–2 competitors	140	22.9%	10.3%	12.6 pp
3 competitors	108	30.6%	9.1%	21.4 pp
4 competitors	190	24.2%	10.1%	14.1 pp
5+ competitors	73	43.8%	12.9%	30.9 pp

Notes: Competitor count is the number of distinct non-Amazon platforms observed for the same product in the analysis panel. Events are isolated upward first moves by Walmart, Target, Best Buy, or Home Depot for which Amazon is observed at the initiating time. The response rate is the share of events followed by an Amazon price increase within 48 hours. The baseline rate is Amazon's probability of raising price within 48 hours in the same competitor-count bin when there is no clean upward first move by one of these major rivals. Response lift is the response rate minus the baseline rate.

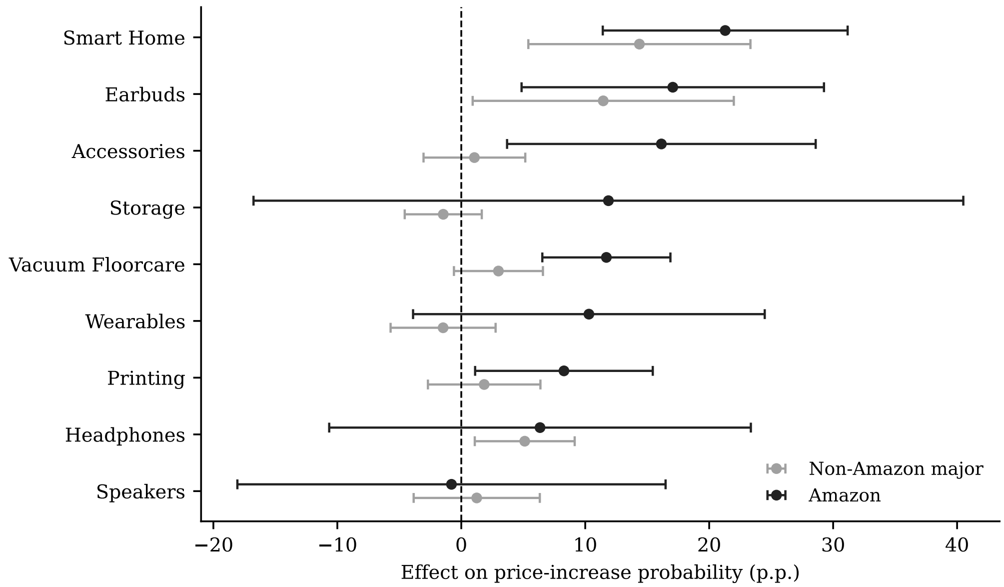

Figure A.1. Amazon Response to Isolated Rival Increases by Product Category.

Notes: The figure reports category-specific versions of the isolated-rival-increase response specification. Black markers report the Amazon coefficient; gray markers report the pooled coefficient for non-Amazon major platforms. Horizontal bars report 95 percent confidence intervals. Each specification includes product, week, day-of-week, and platform fixed effects, with standard errors clustered by product. The figure includes categories with at least 20 Amazon-facing isolated rival increase events and at least five product clusters.

A.2 Empirical State Dependence in Prices

The learning simulations use a parsimonious state: the platform's own lagged price and the lagged lowest rival price. Table A.5 supports this choice. Own prices are highly persistent, and the lagged lowest rival price helps explain the focal platform's current price even after conditioning on its own lagged price. The table also shows that a platform's position relative to the lowest rival price predicts the direction of its next adjustment: platforms priced above the rival benchmark tend to adjust downward, while platforms priced below it tend to adjust upward.

Table A.5. State-Space Motivation: Own and Rival Price States

	(1)	(2)	(3)
	$\log p_{ip,t}$	$\log p_{ip,t}$	$\Delta \log p_{ip,t}$
$\log p_{ip,t-1}$	0.9796***	0.9761***	
	(0.0046)	(0.0058)	
$\log p^{\min}_{-i,p,t-1}$		0.0080***	
		(0.0030)	
$\Delta \log p^{\min}_{-i,p,t-1}$			0.0105
			(0.0070)
$Gap_{ip,t-1}$			-0.0159***
			(0.0045)
Product FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Day-of-week FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Clustered SE	Product	Product	Product
Observations	878,992	878,992	849,328
R^2	0.999	0.999	0.008

Notes: The unit of observation is firm-product-time. The lag $t-1$ is the previous observed state for the same firm-product pair, restricted to observations between 1 and 6 hours apart. $p^{\min}_{-i,p,t}$ is the lowest rival price in the same product market, Δ denotes log price changes, and $Gap_{ip,t-1} = \log p_{ip,t-1} - \log p^{\min}_{-i,p,t-1}$. The sample restricts Amazon and Walmart to first-party listings where seller identity is observed. Standard errors are clustered by product. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Theoretical Appendix

B.1 Comparative Statics for Sequential Matched Prices

Figure A.2 plots the highest sustainable matched price as a function of demand asymmetry and the grid step.

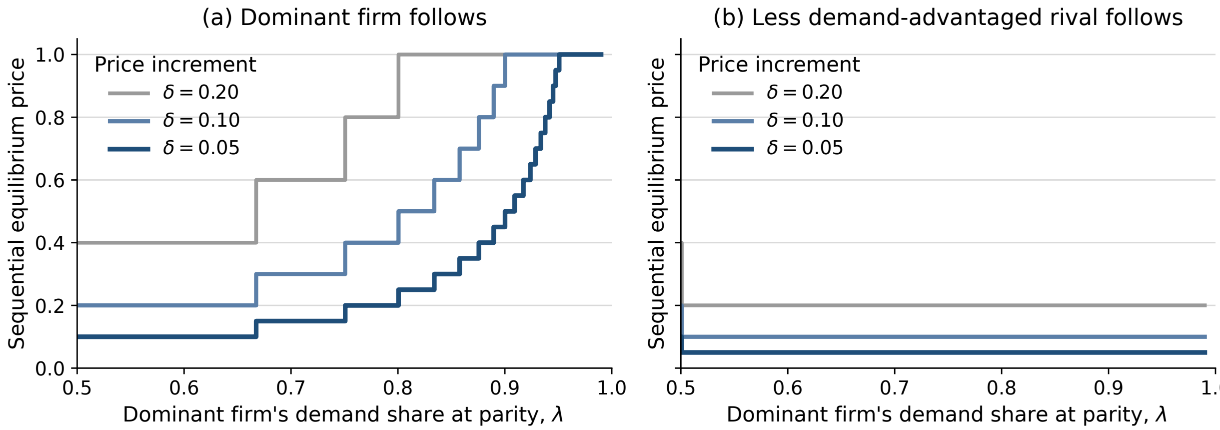

Figure A.2. Sustainable Matched Prices under Sequential Pricing.

Notes: Panel (a) shows the case in which the dominant firm follows, while Panel (b) shows the case in which the rival follows. Prices are the highest sustainable matched prices under the relevant follower constraint. Different lines correspond to different grid steps δ; smaller δ gives a finer price grid, so equilibrium prices adjust in smaller steps.

B.2 Behavioral Separation and Role Preference

Proposition 1 has two implications. First, there is a price region in which the dominant follower prefers matching while the rival follower prefers undercutting, so the same higher price can be sustained only when the dominant firm follows. Second, among matched sequential outcomes, both firms can prefer the role assignment in which the demand-advantaged firm follows.

Corollary 1 (Behavioral separation under demand asymmetry)**.** *Suppose $\lambda > 1/2$. For any first-mover price $p \in \mathcal{P}$ satisfying*

$$\frac{\delta}{\lambda} < p < \frac{\delta}{1-\lambda},$$

the dominant firm matches as the follower, while the rival undercuts.

Let $\Delta_i^U(p)$ denote firm i's net gain, as follower, from undercutting a first-mover price p, relative to matching. For the dominant firm,

$$\Delta_D^U(p) = (1-\lambda)(p-\delta) - \lambda\delta = (1-\lambda)p - \delta.$$

For the rival,

$$\Delta_S^U(p) = \lambda(p-\delta) - (1-\lambda)\delta = \lambda p - \delta.$$

The incentives have opposite signs exactly when

$$\Delta_D^U(p) < 0 < \Delta_S^U(p) \quad \Longleftrightarrow \quad \frac{\delta}{\lambda} < p < \frac{\delta}{1-\lambda}.$$

Corollary 2 (Role preference among matched outcomes)**.** *Among matched sequential outcomes, both firms weakly prefer $S \to D$ to $D \to S$.*

Let $S \to D$ denote the outcome in which the rival moves first and the dominant firm follows, and let $D \to S$ denote the reversed role assignment. Under $S \to D$, the matched price is p_D^{follow}; under $D \to S$, it is p_S^{follow}. At any matched price p, payoffs are

$$\Pi_D(p) = \lambda p, \qquad \Pi_S(p) = (1-\lambda)p.$$

Thus, $\Pi_i^{S\to D} = \Pi_i(p_D^{follow})$ and $\Pi_i^{D\to S} = \Pi_i(p_S^{follow})$. Since Proposition 1 gives $p_D^{follow} \geq p_S^{follow}$, both firms earn weakly higher payoffs under $S \to D$ among matched outcomes.

B.3 Proof of Proposition 2, Part (i)

The game is solved by backward induction. I first solve the deterministic case in which the dominant firm receives the final response opportunity. The calculations verify that, on the normalized grid with $\lambda \in (2/3, 3/4)$ and $\beta > 1/2$, the dominant firm optimally starts at δ, the rival raises to 3δ, and the dominant firm matches. The next appendix subsection solves the reverse deterministic order. I then collect the response-probability extension after both deterministic cases.

Step 1: Solve the final response opportunity. The last move is a one-period pricing problem because only the second market payoff remains. The dominant firm's final response solves

$$a_3^*(a_2) \in \arg\max_{a_3 \in \mathcal{P}} \pi_D(a_3, a_2).$$

Given a rival price p, matching yields λp, while undercutting by one grid step yields $p - \delta$. Thus the dominant firm weakly prefers to match p if and only if

$$\lambda p \geq p - \delta,$$

or $p \leq \delta/(1-\lambda)$. Let

$$p^M \equiv \left\lfloor \frac{\delta}{1-\lambda} \right\rfloor_{\mathcal{P}}$$

denote the highest grid price the rival can induce the dominant firm to match.

Step 2: Solve the second mover's best response. Suppose the dominant firm initially posts δ. Because δ is the lowest positive grid price, the rival cannot gain by undercutting. Its choice is therefore between staying at the low tied price and raising price to induce a future match. Formally, the rival's response solves

$$a_2^*(a_1) \in \arg\max_{a_2 \in \mathcal{P}} \left\{ \pi_S(a_1, a_2) + \beta \pi_S\big(a_3^*(a_2), a_2\big) \right\}.$$

If the rival ties at δ, then $a_2 = a_1$, so the dominant firm's final turn leaves the tie unchanged. The rival's payoff from tying is therefore $(1+\beta)(1-\lambda)\delta$. If the rival instead raises price to p^M, the dominant firm matches in the final move, giving the rival payoff $\beta(1-\lambda)p^M$. Lower price increases yield a smaller matched continuation price, while prices above p^M are not matched. The rival therefore prefers raising price to p^M whenever

$$\beta(1-\lambda)p^M > (1+\beta)(1-\lambda)\delta \quad \Longleftrightarrow \quad \beta > \bar{\beta} \equiv \frac{\delta}{p^M - \delta}.$$

Step 3: Choose the equilibrium opening price. The threshold $\bar{\beta}$ is conditional on the dominant firm posting δ. It remains to verify that δ is the dominant firm's own best initial choice. Consider $\lambda \in (2/3, 3/4)$, so that $\delta/(1-\lambda)$ lies between 3δ and 4δ and $p^M = 3\delta$. Then $\bar{\beta} = 1/2$, so for $\beta > 1/2$ the rival's best response to an initial dominant price of δ is

$$a_2^*(\delta) = 3\delta.$$

The rival sacrifices current demand to create a higher matched continuation state.

Applying the same response problem to the other possible initial dominant prices, for $\beta > 1/2$ in this region, gives

$$a_2^*(0) = 3\delta, \quad a_2^*(\delta) = 3\delta, \quad a_2^*(2\delta) = \delta, \quad a_2^*(3\delta) = 2\delta, \quad a_2^*(4\delta) = 3\delta, \quad a_2^*(5\delta) = 4\delta.$$

Let $U_i(a_1)$ denote firm i's total discounted payoff when the first mover opens with price a_1 and subsequent firms play their best responses. Given this best-response map, the dominant firm's relevant comparison is between starting low at δ and starting high at 5δ. Starting at δ induces the path

$$\delta \to 3\delta \to 3\delta,$$

which gives the dominant firm payoff

$$U_D(\delta) = \delta + 3\beta\lambda\delta.$$

Starting at 5δ instead induces the path

$$5\delta \to 4\delta \to 3\delta,$$

which gives payoff

$$U_D(5\delta) = 3\beta\delta.$$

Thus

$$U_D(\delta) - U_D(5\delta) = \delta\left[1 - 3\beta(1-\lambda)\right] > 0$$

for $\lambda > 2/3$ and $\beta \leq 1$. The remaining openings satisfy

$$U_D(0) = 3\beta\lambda\delta < U_D(\delta),$$

$$U_D(2\delta) = \beta\lambda\delta < U_D(\delta), \qquad U_D(3\delta) = 2\beta\lambda\delta < U_D(\delta),$$

and

$$U_D(4\delta) = 3\beta\lambda\delta < U_D(\delta).$$

Hence, the unique equilibrium path is

$$\delta \to 3\delta \to 3\delta.$$

B.4 Proof of Proposition 2, Part (ii)

This appendix proves part (ii) of Proposition 2. Consider instead the order $S \to D \to M \to S \to M$, so that the rival receives the final revision opportunity. I use the same three backward-induction steps as above.

Step 1: Solve the final response opportunity. The last move again depends only on the second market payoff, but now the rival is the final responder. Its final response is the same follower problem behind p_S^{follow} in Proposition 1: it matches p if and only if $p \leq \delta/\lambda$. Since $\lambda > 1/2$, $\delta/\lambda < 2\delta$, so the rival's matching set never extends above the lowest positive grid price:

$$a_3^*(p) = \begin{cases} 0, & p = 0, \\ \delta, & p = \delta, \\ p - \delta, & p > \delta. \end{cases}$$

Consequently, the price pair $(a_2, a_3^*(a_2))$ at the second clearing is never a tie above δ: it is undercut whenever $a_2 > \delta$, and tied only at δ itself. This holds regardless of what the dominant firm is induced to accept at the first clearing, so no discount factor β sustains a matched price above δ under this order.

Step 2: Solve the second mover's best response. Now fix $\lambda \in (2/3, 3/4)$. If the rival opens at 3δ, the dominant firm matches at the first clearing because matching yields $3\lambda\delta$, while undercutting yields 2δ, and $3\lambda\delta > 2\delta$. But because the rival has the final response opportunity, this match is not preserved: the rival then undercuts to 2δ. Thus opening at 3δ induces the path

$$3\delta \to 3\delta \to 2\delta.$$

If instead the rival opens at 5δ, the dominant firm undercuts to 4δ, and the rival then undercuts back to 3δ. Thus opening at 5δ induces $5\delta \to 4\delta \to 3\delta$.

Step 3: Choose the equilibrium opening price. Using the same payoff notation, the rival's relevant comparison is between these two candidate openings. Opening at 3δ yields

$$U_S(3\delta) = 3(1-\lambda)\delta + 2\beta\delta,$$

while opening at 5δ yields

$$U_S(5\delta) = 3\beta\delta.$$

Opening at 3δ is better exactly when $\beta < 3(1-\lambda)$. The remaining openings are worse. For openings below 3δ,

$$U_S(0) \leq \beta(1-\lambda)\delta < U_S(3\delta),$$

$$U_S(\delta) \leq (1+\beta)(1-\lambda)\delta < U_S(3\delta),$$

and

$$U_S(2\delta) \leq 2(1-\lambda)\delta + \beta\delta < U_S(3\delta).$$

Opening at 4δ induces the dominant firm to undercut to 3δ, after which the rival undercuts to 2δ, so

$$U_S(4\delta) = 2\beta\delta < U_S(3\delta).$$

Opening at 5δ is the only remaining high-opening comparison, and it is worse than 3δ exactly when $\beta < 3(1-\lambda)$. Hence, for $\lambda \in (2/3, 3/4)$ and $\beta < 3(1-\lambda)$, the equilibrium path is

$$3\delta \to 3\delta \to 2\delta.$$

For other parameter values, the exact opening price can differ. The robust point is unchanged: when the rival has the final response opportunity, the second clearing never sustains a matched price above δ.

B.5 Response Probability in the Two-Clearing Game

The deterministic cases above show why the identity of the final responder matters. This extension asks what happens when the dominant firm's final response arrives only with probability ρ_D. If the rival ties at the initial price, $a_2 = a_1$, the dominant firm's final turn is ordinary and ρ_D does not enter. If the rival raises price, $a_2 \neq a_1$, the dominant firm's final move is a response window: with probability ρ_D it can choose a_3, and otherwise its price remains at a_1. I treat ρ_D as an exogenous response-probability parameter, not as a decision variable.

With this modification, initiating p^M gives the rival expected payoff $\beta\rho_D(1-\lambda)p^M$, while matching δ still gives $(1+\beta)(1-\lambda)\delta$. Therefore the rival initiates if

$$\beta\rho_D(1-\lambda)p^M > (1+\beta)(1-\lambda)\delta,$$

or equivalently,

$$\beta > \bar{\beta}_\rho \equiv \frac{\delta}{\rho_D\, p^M - \delta},$$

provided $\rho_D p^M > \delta$. If $\rho_D p^M \leq \delta$, no $\beta \in (0,1]$ rationalizes initiation. Thus, demand advantage determines how valuable the future match can be, while the dominant firm's response probability determines how safe it is for the rival to bet on receiving that match. Appendix Figure A.3 plots this extension: holding λ fixed, the rival's initiation threshold falls as the dominant firm's response becomes more likely.

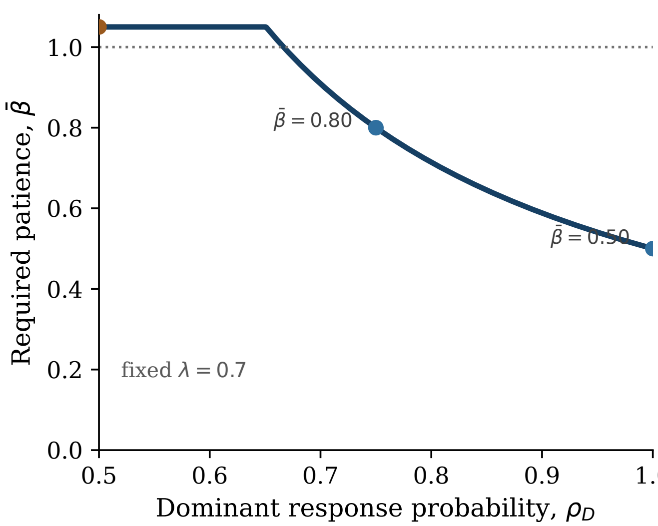

Figure A.3. Response Probability and Costly Rival Initiation.

Notes: The figure fixes $\lambda = 0.7$, so $p^M = 3\delta$, and plots $\bar{\beta}_\rho = \delta/(\rho_D\, p^M - \delta)$, the threshold above which the rival raises price to p^M. Thresholds above one are outside the feasible β range and are shown at the top of the plot.

B.6 Proof of Proposition 3

Fix the candidate policy in Section 3.4 and first consider the benchmark with alternating pricing turns and no stochastic response friction. Let $V_i(p_D, p_S, m)$ denote firm i's value when the posted prices are (p_D, p_S) and firm $m \in \{D, S\}$ moves next. Under the candidate policy, the Bellman equations are

$$
\begin{aligned}
V_D(L,L,D) &= \lambda L + \beta V_D(L,L,S), & V_S(L,L,D) &= (1-\lambda)L + \beta V_S(L,L,S),\\
V_D(L,L,S) &= L + \beta V_D(L,H,D), & V_S(L,L,S) &= \beta V_S(L,H,D),\\
V_D(L,H,D) &= \lambda H + \beta V_D(H,H,S), & V_S(L,H,D) &= (1-\lambda)H + \beta V_S(H,H,S),\\
V_D(L,H,S) &= L + \beta V_D(L,H,D), & V_S(L,H,S) &= \beta V_S(L,H,D),\\
V_D(H,H,D) &= \lambda H + \beta V_D(H,H,S), & V_S(H,H,D) &= (1-\lambda)H + \beta V_S(H,H,S),\\
V_D(H,H,S) &= \lambda H + \beta V_D(H,H,D), & V_S(H,H,S) &= (1-\lambda)H + \beta V_S(H,H,D),\\
V_D(H,L,D) &= \lambda L + \beta V_D(L,L,S), & V_S(H,L,D) &= (1-\lambda)L + \beta V_S(L,L,S),\\
V_D(H,L,S) &= \lambda H + \beta V_D(H,H,D), & V_S(H,L,S) &= (1-\lambda)H + \beta V_S(H,H,D).
\end{aligned}
$$

The one-period-deviation principle then gives eight policy checks, one for each firm whenever it has the pricing turn:

$$
\begin{aligned}
&\lambda L + \beta V_D(L,L,S) \geq \beta V_D(H,L,S), && \beta V_S(L,H,D) \geq (1-\lambda)L + \beta V_S(L,L,D),\\
&\lambda H + \beta V_D(H,H,S) \geq L + \beta V_D(L,H,S), && \beta V_S(L,H,D) \geq (1-\lambda)L + \beta V_S(L,L,D),\\
&\lambda H + \beta V_D(H,H,S) \geq L + \beta V_D(L,H,S), && (1-\lambda)H + \beta V_S(H,H,D) \geq L + \beta V_S(H,L,D),\\
&\lambda L + \beta V_D(L,L,S) \geq \beta V_D(H,L,S), && (1-\lambda)H + \beta V_S(H,H,D) \geq L + \beta V_S(H,L,D).
\end{aligned}
$$

These eight checks reduce to four distinct comparisons: the dominant firm sets L when the rival is low, the rival sets H when the dominant firm is low, the dominant firm sets H when the rival is high, and the rival sets H when the dominant firm is high.

Solving the Bellman equations gives the continuation values needed for these four comparisons:

$$
V_D(H,H,D) = V_D(H,H,S) = \frac{\lambda H}{1-\beta}, \qquad V_S(H,H,D) = V_S(H,H,S) = \frac{(1-\lambda)H}{1-\beta},
$$

$$V_D(L,H,D) = V_D(H,L,S) = \frac{\lambda H}{1-\beta}, \qquad V_S(L,H,D) = V_S(H,L,S) = \frac{(1-\lambda)H}{1-\beta},$$

$$V_D(L,L,S) = L + \beta\frac{\lambda H}{1-\beta}, \qquad V_S(L,L,S) = \beta\frac{(1-\lambda)H}{1-\beta},$$

and

$$V_D(L,H,S) = L + \beta\frac{\lambda H}{1-\beta}, \qquad V_S(H,L,D) = (1-\lambda)L + \beta^2\frac{(1-\lambda)H}{1-\beta}.$$

Substituting these values into the four distinct checks yields the following conditions.

1. The dominant firm's no-self-initiation check is

$$\lambda L + \beta\left(L + \beta\frac{\lambda H}{1-\beta}\right) \geq \beta\frac{\lambda H}{1-\beta},$$

   which is equivalent to

$$\frac{\beta\lambda}{\lambda+\beta} \leq \frac{L}{H}.$$

   This is the dominant firm's waiting condition: if it waits, it earns current low-price demand and matches later; if it initiates, it loses current demand and must wait for the rival to restore the high price.

2. The rival's initiation check is

$$\beta\frac{(1-\lambda)H}{1-\beta} \geq (1-\lambda)L + \beta\left((1-\lambda)L + \beta^2\frac{(1-\lambda)H}{1-\beta}\right),$$

   which is equivalent to

$$\frac{L}{H} \leq \beta.$$

   This is the rival's initiation condition: if it raises now, it loses current demand but gains after the dominant firm matches; if it waits, the low-price state continues.

3. The dominant firm's matching check is

$$\frac{\lambda H}{1-\beta} \geq L + \beta\left(L + \beta\frac{\lambda H}{1-\beta}\right),$$

which is equivalent to

$$\frac{L}{H} \leq \lambda.$$

This is the dominant firm's matching condition: if it matches, it monetizes its parity demand at the high price; if it undercuts, it wins the market at the low price.

4. The rival's no-undercutting check is

$$\frac{(1-\lambda)H}{1-\beta} \geq L + \beta(1-\lambda)L + \beta^3 \frac{(1-\lambda)H}{1-\beta},$$

which is equivalent to

$$\frac{L}{H} \leq \frac{(1-\lambda)(1+\beta+\beta^2)}{1+\beta(1-\lambda)}.$$

This is the rival's persistence condition: if it stays at H, it preserves the high-price path; if it undercuts, it gains current demand but sends the game back through the low-price path.

Combining the four inequalities gives the region stated in Proposition 3.

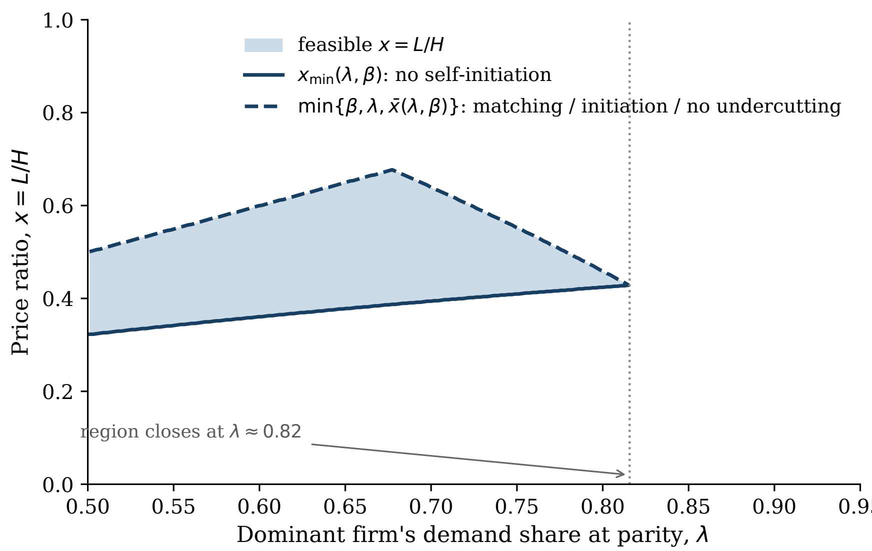

Figure A.4. Sequential Timing Alone Can Support Dominant Followership.

Notes: The figure fixes $\beta = 0.9$ and plots the range of $x = L/H$ for which the candidate policy is a Markov-perfect equilibrium under alternating pricing turns without stochastic response friction, verified by directly solving the Bellman system. The lower curve is the no-self-initiation threshold; the upper curve is the binding minimum of the dominant-matching, rival-initiation, and no-undercutting thresholds.

B.7 Proof of Proposition 4

Fix the candidate policy in Section 3.4, and now allow stochastic response windows. The state is (p_D, p_S, m, c), where $(p_D, p_S) \in \{L, H\}^2$, $m \in \{D, S\}$ is the firm whose pricing turn is next, and $c \in \{0, 1\}$ records whether the previous firm changed price. If $c = 0$, the next pricing turn arrives for sure. If $c = 1$, the next firm receives its response opportunity with probability ρ_i; otherwise its price remains unchanged for that period. Throughout the proposition, $\rho_D = 1$ and $\rho_S \in [0, 1]$.

For firm i, the value of state (p_D, p_S, m, c) satisfies

$$V_i(p_D, p_S, m, c) = \mathbb{E}\left[\pi_i(\tilde{p}_D, \tilde{p}_S) + \beta V_i(\tilde{p}_D, \tilde{p}_S, m', \tilde{c})\right],$$

where $(\tilde{p}_D, \tilde{p}_S)$ is the posted-price pair after the active firm's pricing turn or failed response window, m' is the next scheduled mover, and $\tilde{c}$ records whether the active firm changed price. The expectation is over the response-window realization. Because the policy is fixed, these Bellman equations are linear.

For example, at the off-path state $(H, L, S, 1)$, the rival receives its response opportunity with probability ρ_S. Under the candidate policy, it sets H if the opportunity arrives and remains at L otherwise, so

$$V_i(H, L, S, 1) = \rho_S \{\pi_i(H, H) + \beta V_i(H, H, D, 1)\} + (1 - \rho_S)\{\pi_i(H, L) + \beta V_i(H, L, D, 0)\}.$$

This is where the response probability enters the continuation values used below.

As in the no-friction proof, there are four distinct one-period-deviation checks. The first two are unchanged by ρ_S; the last two pass through the off-path state (H, L), so they are affected by rival response friction.

1. The rival's initiation check is

   $$\beta V_S(L, H, D, 1) \geq (1 - \lambda)L + \beta V_S(L, L, D, 0),$$

   which reduces to $x \leq \beta$. The rival sacrifices current demand when it raises first, so initiation requires enough value from the future matched price.

2. The dominant firm's matching check is

$$\lambda H + \beta V_D(H, H, S, 1) \geq L + \beta V_D(L, H, S, 0),$$

   which reduces to $x \leq \lambda$. Matching monetizes the dominant firm's parity demand at H, while undercutting wins the market at L.

3. The dominant firm's no-self-initiation check is

$$\lambda L + \beta V_D(L, L, S, 0) \geq \beta V_D(H, L, S, 1).$$

   Let $\underline{x}(\lambda, \beta, \rho_S)$ be the smallest $x \in [0, 1]$ for which this inequality holds, with $\underline{x} = 0$ if the inequality holds for all $x \in [0, 1]$. This is the waiting condition: if the dominant firm waits, it earns current low-price demand and lets the rival bear the initiation cost; if it raises first, it must wait for the rival to restore the high-price path. When $\rho_S = 1$, this check reduces to

$$\frac{\beta\lambda}{\lambda + \beta} \leq x,$$

   the lower bound in Proposition 3. Lower ρ_S makes restoration by the rival slower, so the payoff from dominant self-initiation falls and the lower bound weakly relaxes.

4. The rival's no-undercutting check is

$$(1 - \lambda)H + \beta V_S(H, H, D, 0) \geq L + \beta V_S(H, L, D, 1).$$

   Solving the Bellman system gives the upper bound

$$x \leq \bar{x}(\lambda, \beta, \rho_S) \equiv \frac{(1 - \lambda)\left[1 - \beta^3\rho_S - \beta^5(1 - \rho_S)\right]}{(1 - \beta)\left[1 + (1 - \lambda)\big(\beta + (1 + \beta)\beta^2(1 - \rho_S)\big)\right]}.$$

   This is the persistence condition: if the rival stays at H, it preserves the high-price path; if it undercuts, it gains current demand but must rebuild the high-price path through a slower response process. When $\rho_S = 1$, the bound becomes

$$x \leq \frac{(1-\lambda)(1+\beta+\beta^2)}{1+\beta(1-\lambda)},$$

the no-undercutting bound in Proposition 3. Lower ρ_S weakly raises $\bar{x}$, because undercutting becomes less attractive when the rival cannot quickly restore the high matched price.

Combining the four checks gives

$$\underline{x}(\lambda, \beta, \rho_S) \leq x \leq \min\{\beta, \lambda, \bar{x}(\lambda, \beta, \rho_S)\}.$$

Starting from (L, L), once the rival has a pricing turn it chooses H, the dominant firm matches, and the market converges to (H, H).

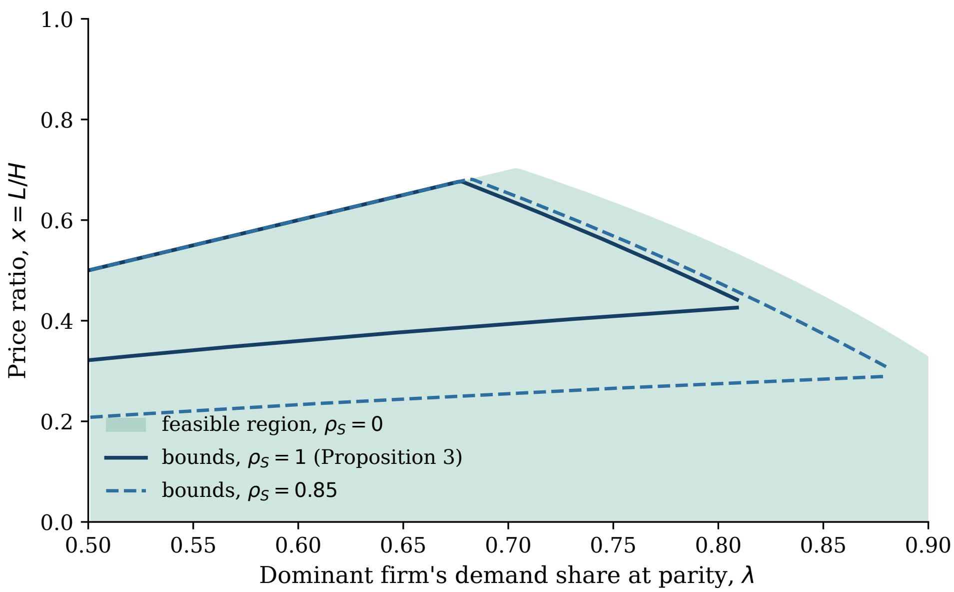

Figure A.5. Response Friction Widens the Feasible Region.

Notes: The figure fixes $\beta = 0.9$ and $\rho_D = 1$ and plots the range of $x = L/H$ sustaining the candidate policy as a Markov-perfect equilibrium, verified by directly solving the Bellman system. The shaded region is the widest case, $\rho_S = 0$; the solid and dashed lines are the bounds at $\rho_S = 1$ (Proposition 3) and $\rho_S = 0.85$, respectively. Lower ρ_S weakly lowers the lower bound and raises the upper bound, so each region contains those for higher ρ_S.

B.8 Loyal-Demand Extension

The baseline model isolates the matching-versus-undercutting mechanism by assuming that a firm charging above its rival receives no demand. This extension relaxes that assumption by allowing each firm to retain loyal consumers even when it charges more than its rival. With loyal demand, the dominant follower can monetize demand advantage in two ways: it can match the rival's price, or it can remain above the rival and serve only loyal consumers. Matching must therefore beat not only undercutting, but also staying expensive.

Proposition 5 (Dominant followership with loyal demand)**.** *Fix a leader price $p = k\delta$, with $k \in \{2, \ldots, M-1\}$. In the loyal-switcher environment, there exists a nonempty open set of primitives (α, β, λ), with*

$$0 < \beta < \alpha, \qquad \alpha + \beta < 1, \qquad \lambda \in [1/2, 1),$$

such that the dominant follower's unique best response is to match p, while the rival's unique best response is to undercut p by one grid step. Thus, the dominant-followership response pattern can arise even when both firms retain positive demand above the leader's price.

Figure A.6 shows that this can hold on a nonempty parameter region. Loyal demand narrows the mechanism, but does not eliminate it: when the dominant firm still gains enough from matching and the rival remains sufficiently dependent on switchers, the behavioral separation behind dominant followership survives.

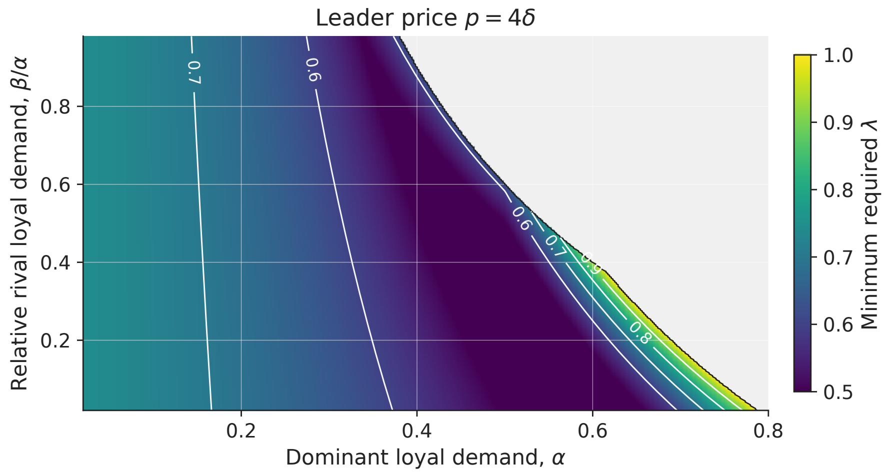

Figure A.6. Dominant Followership with Loyal Demand.

Notes: The figure fixes $M = 5$ and leader price $p = 4\delta$. It plots the minimum λ required for the dominant follower to match while the rival undercuts. Colored regions satisfy $\bar{\lambda}(\alpha, \beta, 4, 5) < 1$; gray regions admit no feasible $\lambda < 1$.

Proof of Proposition 5. Fix a leader price $p = k\delta$, with $k \in \{2, \ldots, M-1\}$. Total demand is normalized to one. A mass α is loyal to the dominant firm, a mass β is loyal to the rival, and the remaining mass

$$s = 1 - \alpha - \beta$$

consists of switchers, with $0 < \beta < \alpha$ and $\alpha + \beta < 1$. Loyal consumers buy from their preferred firm regardless of relative prices. Switchers buy from the lower-price firm, and at equal prices the dominant firm receives share $\lambda \in [1/2, 1]$ of switchers. Demand for the dominant firm is therefore

$$q_D(p_D, p_S) = \begin{cases} \alpha + s, & p_D < p_S, \\ \alpha + \lambda s, & p_D = p_S, \\ \alpha, & p_D > p_S, \end{cases}$$

while demand for the rival is

$$q_S(p_D, p_S) = \begin{cases} \beta, & p_D < p_S, \\ \beta + (1-\lambda)s, & p_D = p_S, \\ \beta + s, & p_D > p_S. \end{cases}$$

The proof first reduces the follower's response problem to three relevant actions, then derives the response inequalities, and finally shows that the feasible set is nonempty.

Step 1: Relevant follower actions. For either follower, matching gives demand at price parity. Among prices below the leader, $p - \delta$ is optimal because any lower price receives the same demand at a lower margin. Among prices above the leader, $M\delta$ is optimal because the firm receives only loyal demand and profit is increasing in price. Thus, it is enough to compare three actions: match, undercut by one grid step, and choose the loyal-only high-price option.

For the dominant follower, the three payoffs are

$$\pi_D^{match} = k\delta[\alpha + \lambda s], \qquad \pi_D^{under} = (k-1)\delta(1-\beta), \qquad \pi_D^{high} = M\delta\alpha.$$

For the rival, they are

$$\pi_S^{match} = k\delta[\beta + (1-\lambda)s], \qquad \pi_S^{under} = (k-1)\delta(1-\alpha), \qquad \pi_S^{high} = M\delta\beta.$$

Step 2: Response inequalities. The dominant follower matches if matching beats both deviations:

$$\pi_D^{match} > \max\{\pi_D^{under}, \pi_D^{high}\}.$$

Substituting the payoffs gives

$$\lambda > \max\left\{\frac{(k-1)(1-\beta) - k\alpha}{ks}, \frac{\alpha(M-k)}{ks}\right\}.$$

The rival undercuts if undercutting beats both alternatives:

$$\pi_S^{under} > \max\{\pi_S^{match}, \pi_S^{high}\},$$

or equivalently,

$$\lambda > \frac{1-\alpha}{ks}$$

and

$$(k-1)(1-\alpha) > M\beta.$$

The first inequality makes undercutting better than matching; the second makes undercutting better than serving only loyal consumers at the high price.

Step 3: Feasibility. Combining the λ-restrictions gives the cutoff

$$\bar{\lambda}(\alpha, \beta, k, M) \equiv \max\left\{\frac{1}{2}, \frac{(k-1)(1-\beta)-k\alpha}{ks}, \frac{\alpha(M-k)}{ks}, \frac{1-\alpha}{ks}\right\},$$

where $s = 1-\alpha-\beta$. The dominant-followership response pattern holds whenever

$$\lambda > \bar{\lambda}(\alpha, \beta, k, M)$$

and

$$(k-1)(1-\alpha) > M\beta.$$

For fixed (α, β, k, M), an admissible $\lambda \in [1/2, 1)$ therefore exists if $\bar{\lambda}(\alpha, \beta, k, M) < 1$ and the additional condition holds.

These conditions are not vacuous. For any $\alpha \in (0, k/M)$, choosing $\beta > 0$ sufficiently small makes $\bar{\lambda}(\alpha, \beta, k, M) < 1$ and satisfies $(k-1)(1-\alpha) > M\beta$. Hence there exists $\lambda \in (\bar{\lambda}(\alpha, \beta, k, M), 1)$. Since the inequalities are strict, the response pattern holds on an open set of primitives.

C Simulation Appendix

C.1 Q-Learning Implementation Details

This appendix collects the exact specification of the Q-learning simulations summarized in Section 4.

Response opportunities. Formally, suppose firm i is active in period t, so firm j is scheduled next. Let ρ_j denote the probability that firm j can revise on its next scheduled turn after firm i changes price. Draw $u_{t+1} \sim U[0,1]$ and define

$$R_{j,t+1} = \mathbf{1}\{a_{i,t} \neq p_{i,t-1}\}\mathbf{1}\{u_{t+1} \leq \rho_j\}.$$

This implements the rule illustrated in Figure 7: $p_{j,t+1} = a_{j,t+1}$ and $p_{i,t+1} = p_{i,t}$ when $R_{j,t+1} = 1$, and prices remain unchanged otherwise. The main simulations vary ρ_D and ρ_S to distinguish symmetric from asymmetric response timing.

Exploration. The ϵ_t-greedy exploration rule is given in Algorithm 1, steps 2–7.

General learning target. Section 4 gives the Q-learning update rule and the simultaneous and two-turn posted-price targets. More generally, a scheduled response window may not be used, in which case the market still clears at unchanged posted prices and the firm earns another realized profit before its next pricing decision. Letting $m_i(t)$ denote the number of market clearings between firm i's choice at t and that next pricing decision, the implemented target is

$$\widehat{Q}^{rw}_{i,t}(s_{i,t}, a_{i,t}) = \sum_{k=0}^{m_i(t)-1} \beta^k \pi_i(p_{i,t+k}, p_{j,t+k}) + \beta^{m_i(t)} \max_{a' \in \mathcal{P}} Q_{i,t}(s_{i,t+m_i(t)}, a').$$

This reduces to the two-turn target in Section 4 when $m_i(t) = 2$ and every scheduled window is used.

Parameters. The discount factor is fixed at $\beta = 0.99$. Learning rates and exploration rates follow parallel hyperbolic decay schedules over a simulation horizon of T periods:

$$\alpha_t = \max\left\{\underline{\alpha}, \alpha_0 \frac{\kappa_\alpha}{\kappa_\alpha + t}\right\}, \qquad \epsilon_t = \max\left\{\underline{\epsilon}, \epsilon_0 \frac{\kappa_\epsilon}{\kappa_\epsilon + t}\right\}.$$

In the baseline simulations, $\alpha_0 = \epsilon_0 = 0.1$, $\underline{\alpha} = \underline{\epsilon} = 0.05$, and $\kappa_\alpha = \kappa_\epsilon = 0.1T$, so both learning and exploration decline gradually over the simulation horizon but remain bounded away from zero.

Myopic best response. Formally, the myopic best response used in the one-learner benchmarks (Section 4) is $BR_i(s_{i,t}) \in \arg\max_{p\in\mathcal{P}} \pi_i(p, p_{j,t-1})$, with ties between matching and undercutting broken toward matching, consistent with the weak inequality in equation (1).

C.2 No-Response Learning Benchmark

Figure A.7 plots the transaction-price trajectories for the simultaneous-updating benchmark discussed in the main text. With no response channel, demand asymmetry does not produce higher learned prices.

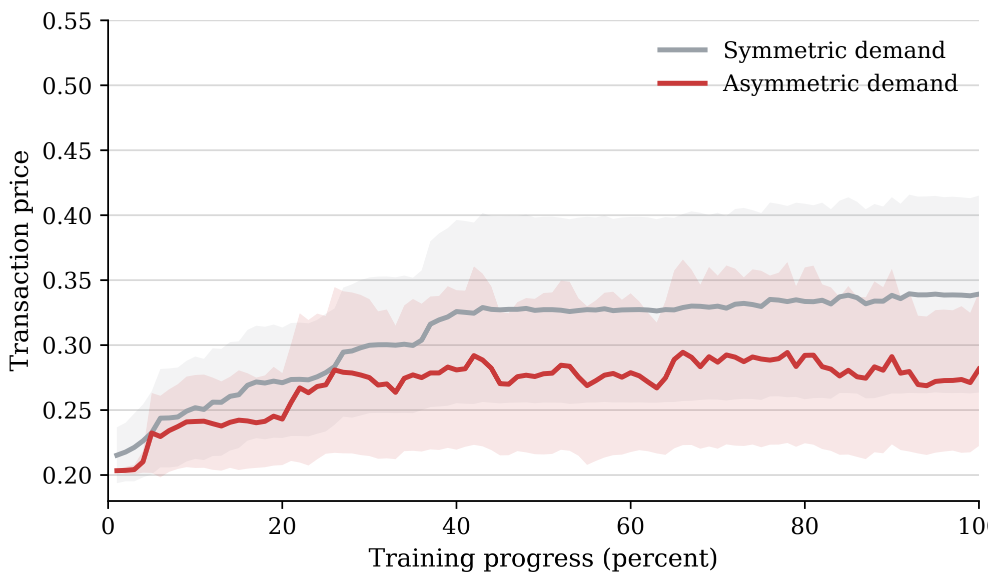

Figure A.7. Transaction Prices under Simultaneous Updating.

Notes: The figure plots average transaction prices in the simultaneous-updating benchmark. Lines average across simulation runs, and shaded regions report 95 percent confidence intervals.

C.3 Late-Period Dynamics in Two-Learner Simulations

The two-learner simulation under dominant followership generates recurrent price adjustment rather than a single fixed price pair. This is the object summarized in Table 9 and Figure 8: late-period play under mutual adaptation.

The intermediate treatments in Figure 8 fall between the symmetric-demand, symmetric-response benchmark and the dominant-followership treatment, matching the same ranking as Table 9: the isolated demand effect leaves most mass concentrated near the lowest price pair, while the isolated response-timing effect shifts mass into mid-to-high price cells.

Learning nevertheless moves play toward a stable region even though the exact price pair does not settle. Figure A.8 shows that the share of periods in which both firms post prices weakly above 0.6 rises during training and is highest when the low-friction responder is also demand advantaged: play becomes concentrated in the high-price region even though the specific price pair within it continues to change.

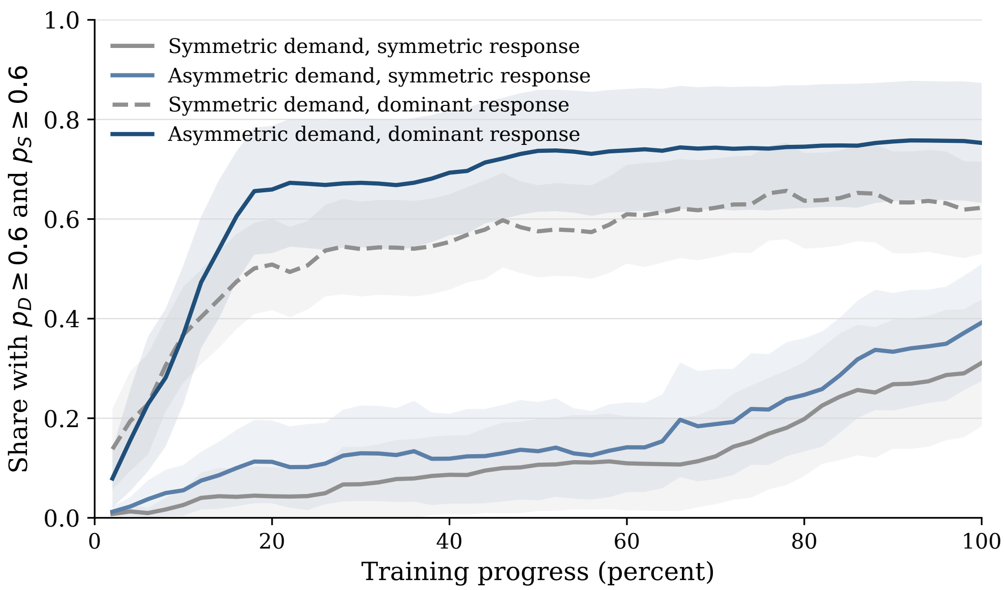

Figure A.8. Learning into High-Price States.

Notes: The figure reports, by training bin, the share of periods in which both firms post prices weakly above 0.6. Lines average across simulation runs; shaded regions report 95 percent confidence intervals across runs.

This interpretation is consistent with prior work on multi-agent Q-learning, where learned play often takes the form of recurrent price dynamics rather than a fixed price profile (Klein 2021; Banchio and Mantegazza 2023; Asker, Fershtman, and Pakes 2024). The one-learner decomposition in Table 13 provides the cleaner evidence on the economic channel behind this pattern.

C.4 Response Decomposition

Figure A.9 decomposes firm D's next-period action after the rival posts above its inherited price. The main text focuses on matching because matching reinforces the rival's higher-price move. This figure shows what happens when matching does not occur. In the comparison cells, higher-price moves are often followed by undercutting or no response. When demand advantage and asymmetric response timing are combined, matching accounts for most of firm D's next-period actions and undercutting becomes much less frequent.

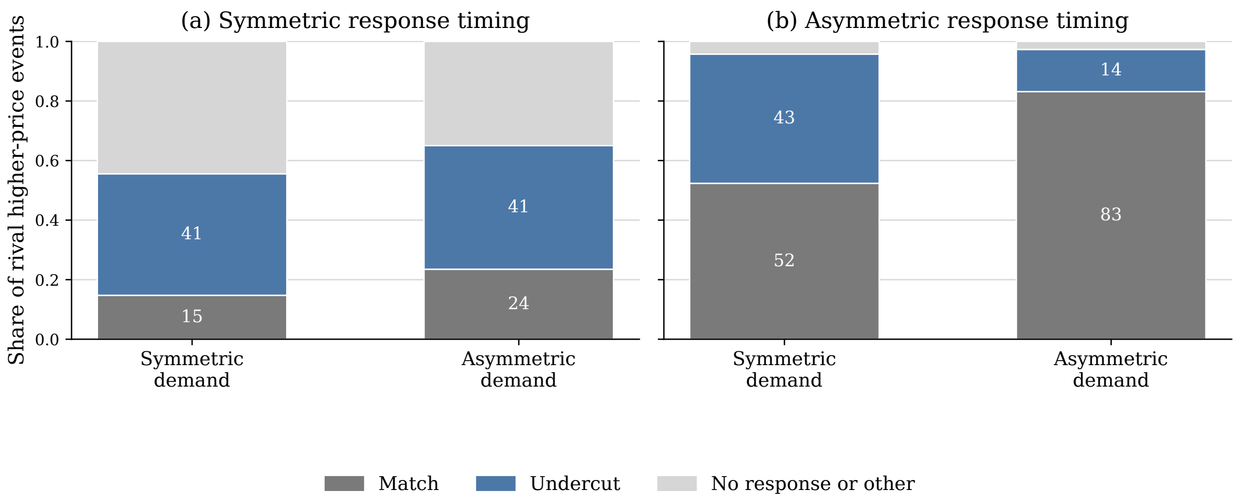

Figure A.9. Response Decomposition after Rival Higher Prices.

Notes: The figure conditions on final-5-percent periods in which firm S posts a price above firm D's inherited posted price. Bars classify firm D's next-period action. A match means firm D sets the same price as firm S; an undercut means it sets a lower price. The remaining category includes cases in which firm D's next posted price is unchanged or is above firm S's price.

C.5 Asymmetric Incidence across Runs

Figure A.10 reports the run-level counterpart to the asymmetric incidence documented in Section 4: it plots, for each simulation run, the rival's posted-price premium against the dominant firm's profit premium. Points in the upper-right region therefore correspond to runs in which the rival posts higher prices on average and the dominant firm earns higher profits. The dominant-followership runs are concentrated in this region, showing that the asymmetric incidence is systematic across runs and not an artifact of period averaging.

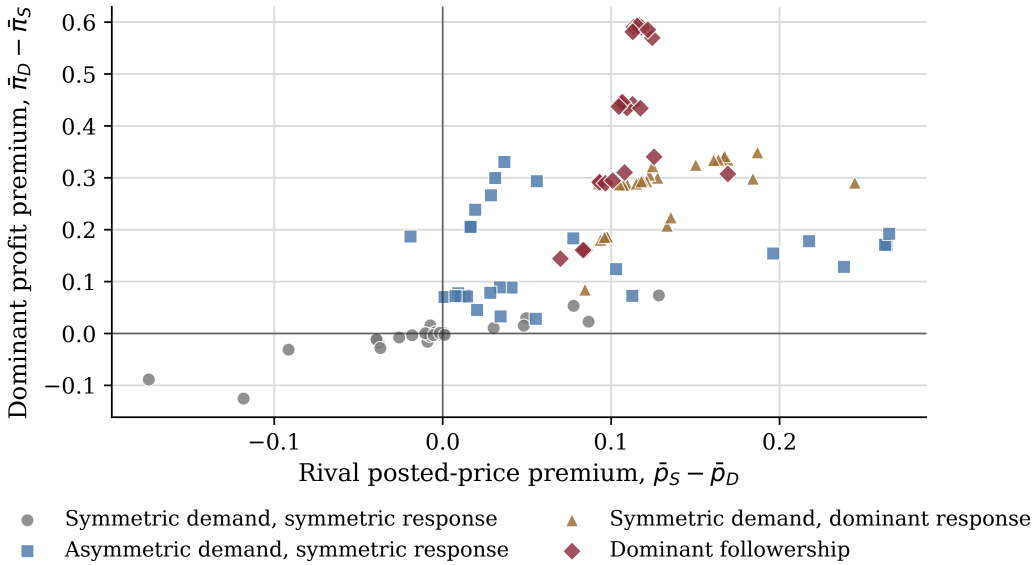

Figure A.10. Run-Level Incidence of Prices and Profits.

Notes: Each point is one simulation run, using averages over the final 5 percent of periods. The horizontal axis reports $\bar{p}_S - \bar{p}_D$, the difference between the rival's average posted price and the dominant firm's average posted price. The vertical axis reports $\bar{\pi}_D - \bar{\pi}_S$, the corresponding dominant-firm profit premium.