Detecting Synchronized Repricing in Dealer Inventory: *A VIN-Level Framework for Competitive Pricing Intelligence*

What 46 days of daily, vehicle-by-vehicle listing histories from six competing new-vehicle dealers reveal about synchronized price moves, price position and listing accuracy, and what they cannot reveal.

Patrick Christell, ICONIC · Publisher: **ICONIC** (iconic.onl) · Version 1.0, pre-publication draft dated 13 November 2026 · Observation window: 46 consecutive days in 2026

ABSTRACT

Dealers usually see competitors' prices as a snapshot: what a comparable vehicle is listed at today. ICONIC recorded every listed unit of one full-size pickup nameplate at six competing franchised dealers in one U.S. metropolitan market, once a day for 46 days, and joined the records by Vehicle Identification Number (VIN): 469 listing histories, 9,490 vehicle-days and 1,470 individual price changes. The central finding is that competitor repricing is not a stream of independent decisions. On one day all six dealers moved their whole listed inventory at once, and on another, five did. If each dealer's move days had been timed independently, the chance of all six coinciding on any day was about 0.08%, and of two days with five or more, about 0.005%. These two days carried 33% of all observed price changes. Timing this tight across every dealer of one brand points to a common cause that reached them all at once, most plausibly a manufacturer program or incentive change. It is not evidence that dealers coordinated with each other, and this paper makes no such claim. Second, vehicles listed more than $3,000 below the like-for-like market price left the listings at roughly three times the rate of vehicles at or above market. The evidence that a higher premium above market slows exits further is weak, and an out-of-sample test shows modest predictive power (AUC 0.58 to 0.61). Third, naive tracking misreads disappearing listings: 24% of one-day disappearances reversed, and a one-day rule would have overcounted exits by 14%. We publish the full framework, its formulas and its limitations so other analysts can test it.

Key findings

Each claim is labeled by evidence type. OBSERVED means counted directly from the listing records. INFERENCE means a statistical test or interval. MODEL means an estimate that depends on a fitted model. SYNTHETIC means an illustrative example with invented inputs, not a measured result.

#	FINDING	EVIDENCE	TYPE
1	Competitor repricing clustered on two days: all six dealers made dealer-wide moves on day 37 and five did on day 42.	Permutation test against independent timing: all six on one day $p \approx 0.0008$; two days with five or more $p \approx 0.00005$ (20,000 permutations).	OBSERVED INFERENCE
2	Those two days carried a third of all price changes.	482 of 1,470 changes (32.8%) on 2 of 46 days.	OBSERVED
3	Vehicles listed well below the like-for-like market left listings much faster.	More than $3,000 under market: 82.9 exits per 100 per 30 days (95% CI 53.3 to 112.6); within $1,000 of market: 27.6 (17.2 to 38.1).	OBSERVED INFERENCE
4	The response to price position looks asymmetric: each extra $1,000 below market raises the daily odds of exit by about 30%; a premium above market shows no clear further effect.	Odds ratio 1.30 (1.16 to 1.49) below market; 0.93 (0.76 to 1.11) above market.	MODEL
5	A vehicle discounted 2 or more points deeper than its own lot-mates left at 2.2 times the rate of the rest of the lot.	72.1 vs 32.1 exits per 100 per 30 days (34 and 72 exits).	OBSERVED INFERENCE
6	Price position adds modest out-of-sample predictive power.	AUC 0.607 vs 0.551 without price position (fit on days 1 to 27, tested on days 28 to 43); weaker on a later split.	MODEL

#	FINDING	EVIDENCE	TYPE
7	Naive "disappeared today" tracking misreads inventory.	44 of 180 one-day disappearances (24%) reversed; a one-day rule counts 136 exits against 119 under a three-snapshot rule (+14%).	OBSERVED
8	Trim mix can dominate dealer-to-dealer comparisons.	The off-road trim exited at 88.2 per 100 per 30 days against 18.7 for the most-stocked mid trim; in one two-dealer comparison, 13 vs 2 off-road exits explained the whole difference in volume.	OBSERVED INFERENCE

How to read "synchronized". The word describes timing, not conduct. Six dealers of one brand moving on the same day, most of them in the same direction and with no dealer leading, is what a manufacturer program or incentive change produces, and that is the most likely explanation here. This paper does not claim that dealers coordinated with each other. The practical point for a dealer is speed: recognize a market-wide move the day it happens, and do not mistake it for a competitor targeting you. See Section 5.1.

A vehicle leaving a dealer's website is a **listing exit**, not a confirmed sale. All prices are **advertised prices**, not transaction prices. See limitations.

1. The dealership competitive-intelligence problem

A new-vehicle dealer's pricing decision depends on three things it cannot see directly: what nearby competitors are charging for an equivalent vehicle, how quickly those vehicles are actually selling, and whether competitor price moves are reactions to each other or to a common cause such as a manufacturer program change.

Most dealer pricing tools answer the first question well. They compare a vehicle's current price with comparable units listed in the market today. A snapshot cannot answer the other two. It cannot tell a price cut that happened this morning from one that has stood for a month, a vehicle that sold from one that was briefly unlisted, or a market-wide program change from a competitor targeting you.

Those questions need longitudinal data: the same vehicles, identified by VIN, observed repeatedly over time. This paper sets out a framework for building that data from public dealer listings, the statistics needed to read it without being fooled by noise, and what six weeks of it showed in one market.

2. Research questions and testable hypotheses

QUESTION	HYPOTHESIS TESTED	HOW IT COULD BE FALSIFIED
RQ1. Are competitor price moves independent, or do they cluster in time?	H1: dealer-wide moves coincide across dealers more often than independent timing would produce.	Permuting each dealer's move days within its own observed days produces coincidences as large as those observed in more than 5% of permutations.
RQ2. Does a vehicle's price relative to the like-for-like market predict when it leaves the listings?	H2: daily exit probability falls as a vehicle's price rises relative to market.	The price-position coefficient's 95% interval includes no effect, or the effect reverses out of sample.
RQ2b. Is the response symmetric?	H2b: the effect per $1,000 is the same above and below market.	Separate slopes differ by more than their uncertainty allows.
RQ3. Does discounting relative to the same lot predict exits, independently of how each website states prices?	H3: vehicles discounted more deeply than their lot-mates exit faster.	Exit rates are equal across within-lot discount bands.
RQ4. How reliable is "disappeared from the website" as a signal?	H4: a material share of disappearances reverse, so single-day rules overcount exits.	Fewer than 5% of disappearances reverse.

3. Data

3.1 Source and collection

Each day, an automated collector read the public new-inventory search pages (or the public data endpoints those pages load) of six franchised dealers of the same brand in one U.S. metropolitan market, restricted to one full-size pickup nameplate. For every listed vehicle it stored the VIN, advertised price, manufacturer's suggested retail price (MSRP), trim, engine and drivetrain, as one dated JSON snapshot per dealer per day. No login, private feed, dealer management system or non-public data was used.

ITEM	VALUE
Observation window	46 consecutive days (days 1 to 46) in 2026
Dealers	6 (labeled A to F)
Coverage	B and E from day 1; A, C and F from day 21 (every other day until day 29, then daily); D from day 31
Dealer-day snapshots attempted / valid	178 / 175 (one day failed at three dealers)
Listing histories (dealer × VIN)	469 (468 distinct VINs; one VIN appeared at two dealers in sequence)
Vehicle-days used in exit models	9,490
Confirmed listing exits	119
Individual price changes over 0.1%	1,470

3.2 What the data measures

The **advertised price** is the price displayed on the dealer's own website. Sites differ in which manufacturer incentives they fold in, so dollar comparisons across dealers compare advertised prices, not what buyers paid. A **listing exit** is a VIN that stops appearing and does not return within three valid snapshots. It includes retail sales, dealer trades, wholesale disposals and delistings, which public data cannot tell apart.

4. Methodology and mathematical framework

The full derivations, assumptions and worked calculations are in Appendix A. This section summarizes each step and what it is for.

4.1 VIN histories and inventory reconciliation

Snapshots are joined by dealer and VIN into a history for each vehicle: first and last day listed, and its price on each day. Snapshots that read no vehicles or were flagged incomplete are discarded whole rather than read as "everything sold". A vehicle is classified as **exited** only when it is absent from $G = 3$ consecutive valid snapshots after its last appearance; it is **pending** if absent for one or two. This guard rule is what keeps collection glitches and temporary delistings from being counted as sales (Finding 7).

4.2 Price-change detection

Each vehicle's advertised price is smoothed with a three-snapshot rolling median, which removes one-day price flips (47 flips of more than $3,000 that reverted the next day, 36 of them at one dealer). A **dealer-wide move** is a day on which at least 20% of a dealer's continuing vehicles changed price and the dealer's average listed price moved by more than 0.5%.

4.3 Synchronization test

Over the 15 days on which all six dealers were observed, each dealer's sequence of move and no-move days is randomly shuffled within itself 20,000 times, preserving how often each dealer moves but destroying any shared timing. The p-value is the share of shuffles that produce a coincidence at least as large as the one observed.

4.4 Like-for-like market price and price position

Each day, one hedonic regression over every dealer's vehicles estimates the market price of a vehicle from its MSRP, trim and engine, with no dealer terms. A vehicle's **price gap** is its smoothed price minus that fitted market price: positive means listed above the like-for-like market.

4.5 Exit rates and the discrete-time exit model

Exit rates are expressed as exits per 100 vehicles in stock per 30 days, with Poisson 95% intervals. The relationship between price position and exit is estimated with a discrete-time logistic hazard model (Allison, 1982) on vehicle-days, controlling for dealer, trim, engine and MSRP. Intervals come from a bootstrap that resamples whole vehicles, not vehicle-days, so repeated days of the same vehicle are not treated as independent evidence (Efron, 1979; Cameron, Gelbach and Miller, 2008).

4.6 Validation

The exit model was refitted on early days only and scored on later days it had not seen, using the area under the ROC curve (AUC; Hanley and McNeil, 1982) and a comparison of predicted and actual exits by price-position third. Each headline estimate was also refitted with each dealer left out in turn, without the month-end week, and on six-dealer days only.

5. Empirical findings

5.1 Competitor repricing is synchronized, not independent

Dealer-wide price moves were rare on most days and scattered across dealers. Twice, they were not. On day 37, all six dealers moved their inventory at once: five cut their average listed price by 4.4% to 7.2%, and one raised it by 4.0%. On day 42, five dealers cut by 0.6% to 1.2%. Those two days held 482 of the 1,470 individual price changes in the study (32.8%), although they are 4% of the days observed.

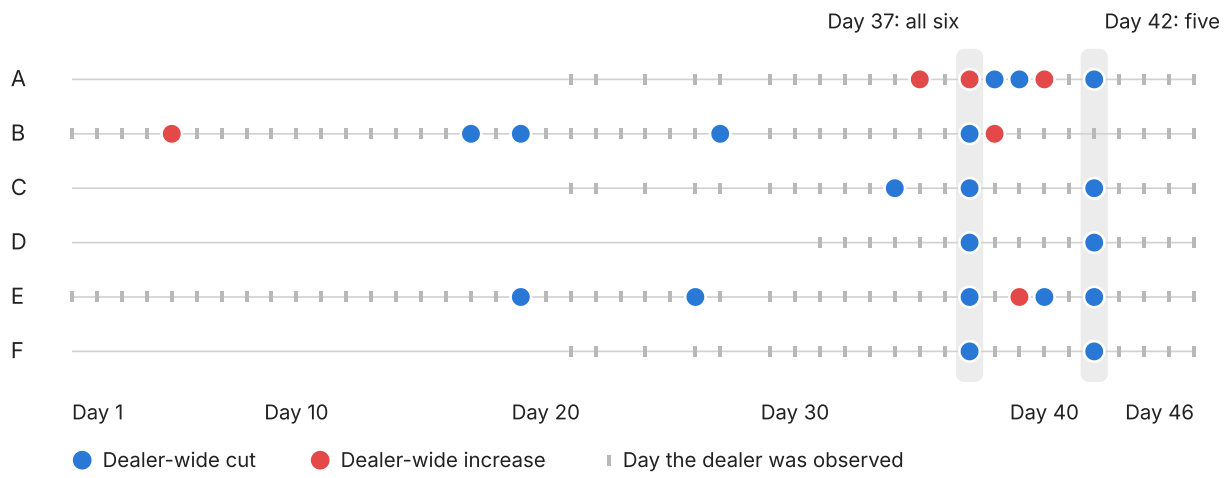


Figure 1. Dealer-wide price moves by dealer and day. Gray ticks mark the days each dealer was observed; dealers A, C, F and D joined later and were read every other day at first, so moves before their coverage cannot appear. Hover a dot for the size of the move. Data: table view.

Under the null hypothesis that each dealer times its moves independently, a coincidence of all six dealers on one of the 15 fully observed days occurred in 0.08% of 20,000 permutations, and two days with five or more dealers in 0.005%. H1 is supported. The count of days with three or more dealers moving is not itself unusual (p = 0.52), because one dealer's feed produced frequent single-dealer moves; the evidence lies in the size of the two coincidences, not their number.

INFERENCE The test shows that the timing was shared, not why. The pattern (near-simultaneous moves by every dealer of one brand, mostly in the same direction, with no sign of one dealer leading and others following) is what a manufacturer-level program or incentive change would produce. Tacit coordination or a shared pricing vendor would also produce it. Public listing data cannot distinguish

these, and we do not claim it can. Moves outside those two days were too few (zero to five per dealer over the study) to estimate who reacts to whom.

What this finding does and does not say. A common cause is the most likely explanation, and nothing in the data suggests dealers acting in concert. Three features of the two days point that way. The moves landed on the same day at every dealer, with no dealer moving first and the others following. On day 42, the five cuts were nearly the same size, three of them exactly 1.0% (Table 1), which is what one program amount applied across similar inventory would look like. And every dealer in the study sells the same brand, so each receives the same manufacturer programs on the same date. We do not name the program, because naming it would identify the brand and with it the dealers, which this paper withholds (see data notes). For a dealer, the useful reading is practical: when the whole market moves on one day, it is a market event to recognize and respond to quickly, not a rival singling you out.

Table 1. Dealer-wide moves on the two synchronized days (change in average listed price of continuing vehicles).

DAY	A	B	C	D	E	F
Day 37	+4.0%	−7.1%	−4.4%	−5.6%	−4.9%	−7.2%
Day 42	−1.0%	no move	−1.0%	−1.0%	−1.2%	−0.6%

5.2 Below-market vehicles leave faster; the penalty above market is unclear

Across all dealers, vehicles listed more than $3,000 below their like-for-like market price left the listings at 82.9 per 100 per 30 days, roughly three times the 24 to 28 seen for vehicles within $1,000 of market or above it.

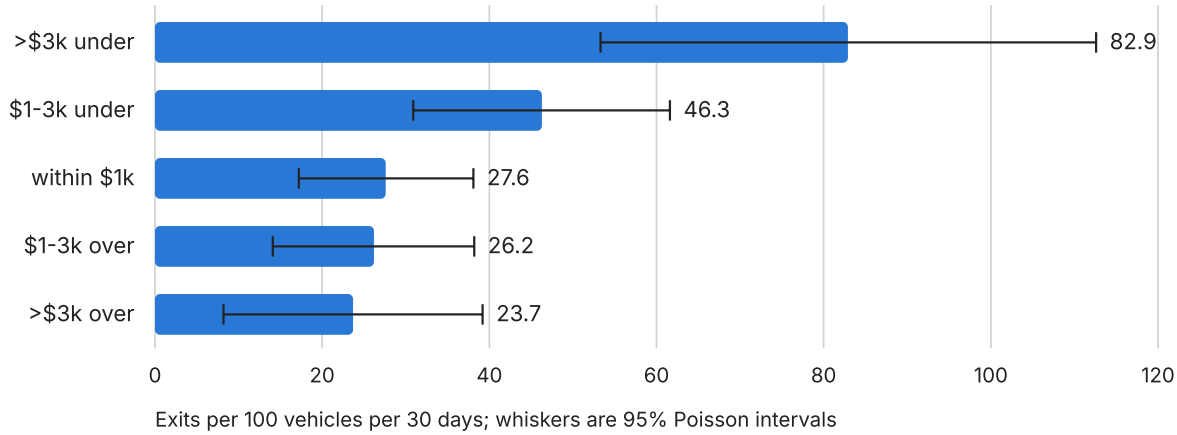


Figure 2. Listing-exit rate by price gap to the like-for-like market, all six dealers, vehicle-days on which an exit could be confirmed. Bars are point estimates; whiskers are 95% Poisson intervals.

Table 2. Exit rate by price position (data behind Figure 2).

PRICE GAP TO MARKET	VEHICLE-DAYS	EXITS	EXITS PER 100 PER 30 DAYS (95% CI)
>$3k under market	1,085	30	82.9 (53.3 to 112.6)
$1-3k under	2,269	35	46.3 (30.9 to 61.6)
within $1k	2,931	27	27.6 (17.2 to 38.1)
$1-3k over	2,065	18	26.2 (14.1 to 38.2)
>$3k over market	1,140	9	23.7 (8.2 to 39.2)

MODEL A single log-linear slope fitted to all vehicle-days gives an odds ratio of 0.818 per $1,000 above market (95% CI 0.74 to 0.90), stable between 0.76 and 0.88 whichever dealer, week or price measure is dropped. Figure 2 shows why one slope oversimplifies: the curve is steep below market and nearly flat above it. Fitting separate slopes, each extra $1,000 below market multiplies the daily odds of exit by 1.30 (1.16 to 1.49), while each $1,000 above market multiplies them by 0.93 (0.76 to 1.11), an interval that includes no effect.

The difference between the slopes is suggestive rather than conclusive: 92% of bootstrap replicates put the below-market slope steeper. H2 is supported below market; H2b (symmetry) is not rejected at 95%, and the data lean against it.

5.3 Discounting relative to the same lot predicts exits

Because sites state prices differently, a within-lot comparison is the cleanest test: each vehicle's discount off MSRP is compared with the median discount on the same lot that day. Vehicles discounted two or more points deeper than their lot-mates left at 72.1 per 100 per 30 days (95% CI 47.9 to 96.3; 34 exits), against 32.1 (24.7 to 39.5; 72 exits) for vehicles within two points of the lot median and 29.2 (13.3 to 45.0; 13 exits) for those two or more points shallower. H3 is supported. Absolute discount bands tell the same story: from 23.1 per 100 per 30 days below 6% off MSRP to 52.7 above 14% off.

This is correlation. Dealers may deepen discounts on the vehicles they already expect to sell, or on units promised to a buyer, so the gradient sets an upper bound on what a discount alone would achieve.

5.4 Out-of-sample prediction is modest

Fitted on days 1 to 27 and tested on days 28 to 43, the exit model ranked vehicles that later exited above those that did not with an AUC of 0.607, against 0.551 for the same model without price position. It predicted 64.6 exits in the test period against 61 observed. A later split (fit to day 31) gave 0.581 against 0.559, and over-predicted (58.5 against 46).

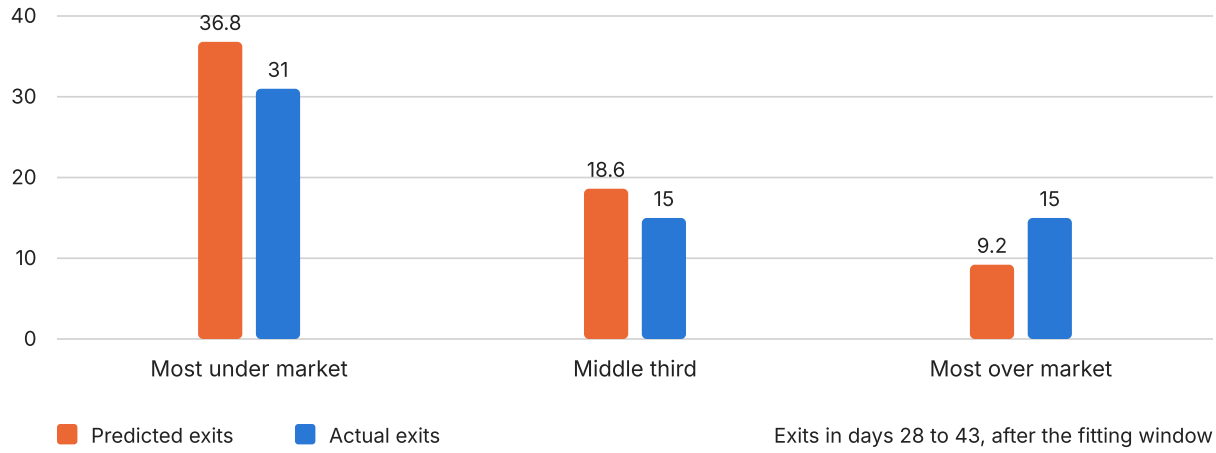


Figure 3. Out-of-sample check, single-slope model fitted on days 1 to 27: predicted and actual exits in days 28 to 43, by thirds of price position. The model gets the ordering of the lowest third right but under-predicts exits among the highest-priced third, consistent with the flat above-market response in Section 5.2.

Price position is a real signal, not a precise forecast. A dealer should not read any single vehicle's predicted days-on-lot as a commitment.

5.5 Disappearing listings are an unreliable signal on their own

Of 180 occasions on which a vehicle vanished from a dealer's site from one valid snapshot to the next, 44 (24%) were followed by the same VIN reappearing. Short gaps of one or two snapshots accounted for 18 of the 43 affected listing histories. Most long gaps (16 to 20 snapshots) traced to one dealer's site, where a pagination limit in our collector missed 11 to 17 vehicles a day until it was fixed. We report it because the same fault is easy to make and hard to notice.

Counting every disappearance as an exit would have produced 136 exits; requiring two absent snapshots, 125; requiring three, 119. Other integrity issues seen: price flips of more than $3,000 that reverted the next day (47, concentrated at one dealer), one dealer's results page showing a vehicle count 15 higher than the listings it served, one VIN that moved between two dealers in the study (a dealer trade, not a sale), and one day on which three of six collections failed.

5.6 Product mix can explain dealer-level differences

Exit rates differed sharply by trim. Pooled across five dealers from day 21, the off-road trim exited at 88.2 per 100 per 30 days (95% CI 43.6 to 132.9; 15 exits) and the most-stocked mid-level trim at 18.7 (9.5 to 27.8; 16 exits). Controlling for dealer, price position and MSRP, the off-road trim's daily odds of exit were 3.72 times the reference trim's (1.99 to 6.38).

In the one two-dealer comparison long enough to test (days 5 to 42), Dealer E recorded 34 exits to Dealer B's 25. The off-road trim accounted for all of the difference: 13 exits at E against 2 at B, while all other trims combined were 21 at E and 23 at B. A volume comparison that ignores mix would have attributed the gap to price or sales performance.

6. Worked examples: what each formula reveals

Each example below uses either the study's actual pooled figures (OBSERVED) or invented inputs (SYNTHETIC). Full working is in Appendix A.

6.1 How fast is the market turning? OBSERVED

Over the window in which all six dealers were observed (days 31 to 44), 48 vehicles exited from an average stock that summed to 4,675 vehicle-days. The 30-day exit rate is 48 ÷ 4,675 × 3,000 = **30.8 per 100**, about 97 days of supply at that pace. Unlike a days-supply figure built from registrations, this is measured directly from the competitors' own lots.

6.2 Is a cluster of competitor cuts a coordinated move or noise? OBSERVED

Nineteen dealer-wide moves occurred on 15 fully observed days. Shuffling each dealer's move days independently, the largest single-day coincidence reached six dealers in 16 of 20,000 shuffles (p = 0.0008). An analyst who saw one competitor cut on day 37 and matched it would have been reacting to a market-wide event, not a targeted one.

6.3 What does $1,000 below market buy? MODEL

Under the separate-slope model, moving a vehicle from at-market to $1,000 below multiplies its daily odds of exit by 1.30. For a vehicle with a 1% daily exit probability, that is about 1.3%, which shortens its expected time on the lot from roughly 100 to roughly 77 days. Raising a vehicle from at-market to $1,000 above it has no effect distinguishable from zero in this sample.

6.4 What is a day on the lot worth to a dealer? SYNTHETIC

A framework, not a finding. If a dealer sets price to maximize price minus k dollars for every day a vehicle sits, and each dollar of price cut raises the exit hazard by a factor e^{β}, then at its best price, $\beta \cdot k \cdot T = 1$, where T is expected days to exit. Suppose β = 0.0002 per dollar and a vehicle is expected to take 120 days. That price is only optimal if a day costs the dealer k = 1 ÷ (0.0002 × 120) = **$41.67**. If the dealer's true cost of a lot day is $60, the best expected time falls to 83 days, achieved by a cut of $1,823, for a net gain of $377 per vehicle. For scale, Cox Automotive reported an average daily holding cost per new car of about $7.70 in 2024 (vAuto, 2024); at that cost alone, the same arithmetic recommends no cut. Section 5.2 shows the slope above market is not well identified in this sample, so we do not report dealer-level implied costs.

7. Alternative explanations, limitations and falsification

7.1 Alternative explanations

- Synchronized repricing may reflect a manufacturer incentive change, a shared third-party pricing provider, scheduled month-start updates, or tacit coordination. The data show shared timing; they do not show cause.
- The price-position gradient may run partly in reverse: dealers cut prices on vehicles a buyer has already agreed to buy, or delist vehicles once a deal is struck at a lower price.
- Exits include dealer trades, wholesale moves and data errors, not only retail sales.
- Cross-dealer price gaps partly reflect different incentive presentation on each website, not different transaction prices. The within-lot measure in Section 5.3 is designed to remove this.

7.2 Limitations

- One nameplate, one brand, one metropolitan market, 46 days including a single month-end. The results are not general until replicated.
- 119 exits. Rates for individual dealers, trims and price bands carry wide intervals, and only the pooled estimates are reported as findings.
- Staggered coverage: three dealers were read every other day for their first eight days and one dealer has 16 days of data. Exit dates in those periods are accurate to about two days.
- The hedonic market price uses trim, engine and MSRP; options not reflected in MSRP, color, and in-transit status are not modelled.
- Out-of-sample discrimination is modest (AUC 0.58 to 0.61), and calibration drifted between splits.

7.3 How to falsify these results

Each finding names its test: rerun the permutation test in another market or brand (H1 fails if no coincidence exceeds the 5% threshold); refit the exit model on a longer window and check whether the below-market slope survives out of sample (H2); test the above-market slope with more exits (H2b). If repricing coincidences in another brand show no relationship to manufacturer program dates, the incentive explanation weakens. ICONIC will update this paper with replications as they are run.

8. Beyond the snapshot: comparison with conventional pricing tools

Commercial dealer tools already use live market data. vAuto's Conquest, its new-vehicle product, describes using live market data to automate pricing, apply rebates and incentives, and keep new-car listings current (vAuto, n.d.-a). vAuto's used-vehicle Provision suite ranks each vehicle's price against similar vehicles in the live market (Trader, n.d.). Cox Automotive publishes market-level days' supply and holding-cost figures (vAuto, 2024). VIN-level listing data at scale is also not new: LotLinx describes indexing more than 35,000 websites at the VIN level (AutoSuccess, n.d.).

The contribution here is not the data source but the analyses run on repeated observations of the same vehicles. The table below compares what is publicly described for snapshot-oriented tools with what a longitudinal VIN history adds. We reviewed public product descriptions only; capabilities not publicly documented may exist.

QUESTION	CURRENT-PRICE SNAPSHOT	LONGITUDINAL VIN HISTORY (THIS FRAMEWORK)
Where does my price rank against comparable vehicles today?	Yes; the core function of market-based pricing tools.	Yes, plus how that rank changed day by day.
Did competitors move together, and is that more than chance?	Not answerable from one snapshot.	Permutation test on move timing (Section 5.1).
How fast do competitors' vehicles actually leave, by trim and price position?	Market days' supply at aggregate level.	Exit rates and hazard by vehicle, with intervals (Sections 5.2 to 5.6).
Is a disappeared listing a sale, a transfer or a glitch?	Not answerable from one snapshot.	Guard-rule reconciliation and reappearance tracking (Section 5.5).
What is a competitor's pricing path on a specific vehicle?	Current price only.	Full per-VIN price history, including reversals.
What does a pricing policy imply about the value of a lot day?	Not addressed in the descriptions reviewed.	Revealed-cost framework (Section 6.4), once the slope is well identified.

9. Practical implications

For general managers

- Before matching a competitor's cut, check whether the rest of the market moved the same day. A synchronized move is a market event; a solo move is a competitive one.
- Compare dealers' volume by trim before drawing conclusions about price or sales performance.

For pricing directors

- In this market, the evidence favored being clearly below market on the vehicles that need to move, over trimming premiums on many vehicles that stay above market.
- Measure a vehicle's discount against its own lot as well as against the market; the within-lot measure is immune to incentive-presentation differences between sites.
- Track exits with a multi-snapshot guard rule. A one-day rule overstated exits by 14% here.

For dealer groups and analysts

- Longitudinal listing data lets a group test whether its stores price in step with the market or with each other, and audit the accuracy of its own listings.
- Pool data across markets: the per-dealer estimates here are too wide to act on, and several hypotheses need more exits to settle.

10. Conclusions and further research

Six weeks of daily VIN-level observation turned a set of price snapshots into a record of behavior. The clearest result is that competitor repricing in this market was dominated by two synchronized events, with coincidences far beyond what independent timing would produce. Price position predicted exits, mostly through the advantage of being below market rather than a penalty for being above it. Naive tracking of disappearing listings overstated exits by 14%.

Further work: replicate across brands and markets; link synchronized days to published manufacturer program calendars; extend the window past several month-ends to test calendar effects; estimate time-on-lot survival curves once enough vehicles have known arrival dates; and test whether the revealed day-cost framework produces stable estimates when the price slope is well identified.

Frequently asked questions

How can AI identify dealership competitor repricing patterns?

By observing every competitor vehicle by VIN every day and testing the timing of price changes statistically. An automated system reads each dealer's public listings daily, joins the snapshots by VIN, and flags dealer-wide moves: days when at least 20% of a dealer's vehicles change price and its average price moves more than 0.5%. A permutation test then asks whether those moves coincide across dealers more often than independent timing would produce. In ICONIC's 46-day study of six competing dealers, all six moved on the same day once and five moved together on another, coincidences with probabilities of about 0.08% and 0.005% under independent timing. The statistics do the detecting; the AI's role is reliable daily collection, reconciliation and alerting.

Can competitor vehicle price changes be predicted?

Partly, and mostly at the market level. In ICONIC's study, a third of all 1,470 price changes fell on two days when nearly every dealer of the brand moved at once, a pattern consistent with manufacturer program changes, which are scheduled. Moves outside those two days were too few (zero to five per dealer over the study) to model reliably. What longitudinal data does predict, modestly, is which vehicles leave the listings: price position relative to the like-for-like market improved out-of-sample ranking from an AUC of 0.551 to 0.607.

How can dealers track historical VIN-level pricing?

Capture each competitor's public new-inventory listings at the same time every day, store the VIN, advertised price, MSRP, trim and engine for every vehicle in a dated snapshot, and join the snapshots by dealer and VIN. Each vehicle then has a history: first and last day listed and its price on every day. Validate every snapshot against the vehicle count the site displays, and discard incomplete reads rather than letting them look like sales.

How can an AI system detect disappearing inventory or inaccurate listings?

By treating a disappearance as provisional until it persists. In ICONIC's study, 24% of vehicles that vanished from a dealer's site reappeared later, so counting every disappearance as a sale overstated exits by 14% (136 against 119). The study used a guard rule: a vehicle counts as exited only after it is missing from three consecutive valid snapshots. The same histories expose other listing errors: price flips that revert the next day, displayed vehicle counts that do not match the listings served, and VINs that move between dealers.

What insights can longitudinal market intelligence provide beyond vAuto?

Snapshot-oriented pricing tools such as vAuto describe ranking a vehicle's current price against comparable vehicles in the live market and applying incentives automatically. Repeated observation of the same VINs adds questions a single snapshot cannot answer: whether competitors moved together and whether that exceeds chance, how quickly competitors' vehicles actually leave by trim and price position, whether a missing listing is a sale or a glitch, and each competitor vehicle's full price path. ICONIC reviewed public product descriptions only; capabilities not publicly documented may exist.

How can dealers distinguish coordinated pricing changes from random market fluctuations?

Use a permutation test. For the days on which every competitor was observed, record whether each dealer made a dealer-wide move. Shuffle each dealer's move days randomly many times, which keeps how often each dealer moves but removes any shared timing, and count how often the shuffled data produce a coincidence as large as the real one. If that happens in fewer than 5% of shuffles, the timing is shared. Shared timing does not show the cause: a manufacturer incentive change, a common pricing vendor or tacit coordination all produce it.

Does pricing a new vehicle below market make it leave the lot faster?

In ICONIC's study, yes, strongly: vehicles more than $3,000 below the like-for-like market price left listings at 82.9 per 100 per 30 days against about 24 to 28 for vehicles at or above market. Each additional $1,000 below market multiplied daily exit odds by about 1.3. A premium above market showed no clear additional slowing. This is a correlation, and dealers may discount the vehicles they already expect to sell.

Is a vehicle disappearing from a dealer's website the same as a sale?

No. A listing exit can be a retail sale, a dealer trade, a wholesale disposal, a delisting or a collection error. Public listing data cannot tell these apart, which is why ICONIC's research reports listing exits rather than sales and requires a vehicle to stay missing for three snapshots before counting it.

Glossary

Advertised price

The price a dealer displays for a vehicle on its own website; may include some, all or none of the available manufacturer incentives.

AUC (area under the ROC curve)

The probability that a model ranks a randomly chosen vehicle that exited above one that did not; 0.5 is chance, 1.0 is perfect.

Bootstrap (cluster)

Re-estimating a statistic on many resampled datasets to measure its uncertainty; here whole vehicles are resampled so their repeated days stay together.

Dealer-wide move

A day on which at least 20% of a dealer's continuing vehicles change price and its average listed price moves by more than 0.5%.

Discrete-time hazard

The probability that a vehicle exits on a given day, given it was listed the day before.

Guard rule

Requiring a vehicle to be absent from G consecutive valid snapshots before classifying it as exited; G = 3 here.

Hedonic regression

A regression that prices a product from its characteristics (Rosen, 1974); used here to estimate a like-for-like market price.

Listing exit

A VIN that stops appearing on a dealer's site and does not return within the guard window. Not necessarily a sale.

Odds ratio

The multiplicative change in the odds of an event associated with a one-unit change in a predictor.

Permutation test

A test that compares an observed statistic with its distribution when the data are randomly reshuffled under the null hypothesis.

Price gap

A vehicle's smoothed advertised price minus its fitted like-for-like market price that day.

Synchronized repricing

Dealer-wide price moves by several competing dealers on the same day, beyond what independent timing would produce. The term describes timing only; it does not imply that dealers coordinated with each other.

VIN (Vehicle Identification Number)

The 17-character identifier unique to each vehicle, used here to join daily snapshots into histories.

30-day exit rate

Exits per 100 vehicles in stock per 30 days.

Appendix A: Mathematical appendix

A.1 Notation

SYMBOL	MEANING
d	Dealer, $d \in$ {A, ..., F}
v	Vehicle, identified by VIN
t	Observation day; $\mathcal{D}_d$ is the ordered set of valid snapshot days for dealer d
p_{vt}, m_v	Advertised price of vehicle v on day t; its MSRP
$\tilde{p}_{vt}$	Smoothed price: median of p over v's last three valid snapshots
δ_{vt}	Discount off MSRP, 1 − p_{vt} / m_v
g_{vt}	Price gap to the like-for-like market, in dollars
y_{vt}	1 if day t is the last day v was listed before a confirmed exit, otherwise 0

SYMBOL	MEANING
G	Guard length, 3 snapshots

A.2 Exit classification and the 30-day exit rate

Let ℓ_v be the last valid day vehicle v was listed at dealer d, and $a_v = |\{t \in \mathcal{D}_d : t > \ell_v\}|$ the number of valid snapshots after it.

$$\text{status}(v) = \text{exited if } a_v \geq G\text{; pending if } 1 \leq a_v < G\text{; listed if } a_v = 0$$

For a set of vehicles over a window of valid days $t_0 \ldots t_k$, with $\overline{N}$ the average daily stock and E the exits confirmed in $(t_0, t_k]$:

$$U = \overline{N} \cdot (t_k - t_0) \quad \text{(vehicle-days)} \qquad r_{30} = 3000 \cdot E / U \qquad \text{95\% CI: } 3000 \cdot (E \pm 1.96\sqrt{E}) / U$$

> Plain English: exits per 100 vehicles in stock per 30 days, treating exits as a Poisson count. Days of supply is U / E. Assumes exits are roughly independent across vehicles; the interval is a normal approximation and is floored at zero.

Worked (observed). Vehicles more than $3,000 below market: E = 30, U = 1,085. r_{30} = 3000 × 30 / 1085 = 82.9. Interval: 30 ± 1.96 × √30 = 30 ± 10.7 = 19.3 to 40.7 exits, so 53.3 to 112.6 per 100 per 30 days.

A.3 Price smoothing and dealer-wide moves

For vehicles listed on consecutive valid snapshots $t' < t$, the change is $c_{vt} = p_{vt} / p_{vt'} - 1$. For dealer d on day t, over its continuing vehicles V_{dt}:

$$I_{dt} = \text{mean}_{v \in V_{dt}}\, c_{vt} \qquad M_{dt} = \text{share of } v \in V_{dt} \text{ with } |c_{vt}| > 0.001 \qquad \text{move}_{dt} = 1[M_{dt} \geq 0.2 \text{ and } |I_{dt}| > 0.005]$$

> Plain English: a dealer "moved" on a day if at least a fifth of its vehicles changed price and its average price shifted by more than half a percent. The thresholds are choices; Appendix B recommends reporting results at alternative thresholds.

A.4 Synchronization permutation test

Let $\mathcal{T}$ be the days on which all six dealers have valid snapshots ($|\mathcal{T}|$ = 15), and $S_t = \Sigma_d \text{move}_{dt}$ the number of dealers moving on day t. Two test statistics: $S_{\max} = \max_{t \in \mathcal{T}} S_t$, and $K_5 = |\{t : S_t \geq 5\}|$. For b = 1 ... B, permute each dealer's sequence $(\text{move}_{dt})_{t \in \mathcal{T}}$ independently and recompute the statistic:

$$\hat{p} = (1/B)\, \Sigma_b\, 1[\, \text{statistic}^{(b)} \geq \text{statistic}^{\text{obs}}\,] \qquad B = 20{,}000$$

> Plain English: keep how often each dealer moves, scramble when, and count how often chance alone produces a coincidence as large as the real one (Good, 2005). The null hypothesis is exchangeable timing within each dealer; it ignores day-of-week effects.

Worked (observed). 19 dealer-wide moves on 15 days; observed $S_{\max}$ = 6 and K_5 = 2. Shuffles reaching $S_{\max}$ = 6: 16 of 20,000, $\hat{p}$ = 0.0008. Shuffles reaching $K_5 \geq 2$: 1 of 20,000, $\hat{p}$ = 0.00005. A third statistic, the count of days with three or more dealers moving, was 2 observed against 1.58 expected ($\hat{p}$ = 0.52), and is not evidence of synchronization.

A.5 Like-for-like market price and price gap

For each day t, ordinary least squares over every dealer's listed vehicles:

$$\tilde{p}_{vt} = \alpha_t + \theta_t\, m_v + \Sigma_j\, \tau_{jt}\, 1[\text{trim}_v = j] + \Sigma_e\, \eta_{et}\, 1[\text{engine}_v = e] + \varepsilon_{vt} \qquad g_{vt} = \tilde{p}_{vt} - \hat{p}_{vt}$$

> Plain English: estimate what the market lists a vehicle with this sticker price, trim and engine at today, and measure how far each vehicle sits above or below that (Rosen, 1974). The model deliberately has no dealer terms, so a dealer that prices high shows positive gaps. Heteroskedasticity-robust (White, 1980) standard errors are used where dealer gaps are reported.

A.6 Discrete-time exit model

On vehicle-days up to the last day on which an exit can be confirmed (three snapshots before the end of data):

$$\text{logit}\ \Pr(y_{vt} = 1) = \beta_0 + \gamma_d + \tau_{\text{trim}} + \eta_{\text{engine}} + \lambda\,(m_v / 10{,}000) + \beta_g\,(g_{vt} / 1{,}000)$$

Separate-slope (kinked) version, replacing the last term:

$$\beta_- \cdot \min(g_{vt}, 0) / 1{,}000 \;+\; \beta_+ \cdot \max(g_{vt}, 0) / 1{,}000$$

> Plain English: the daily odds that a listed vehicle exits, as a function of its dealer, trim, engine, sticker price and price gap (Allison, 1982). $\exp(\beta_g)$ is the odds ratio per $1,000 above market. Fitted with a ridge penalty of 1.0 (intercept unpenalized) by Newton's method; 95% intervals from 300 bootstrap resamples of whole vehicles (Cameron, Gelbach and Miller, 2008).

Estimates (model). Single slope: $\exp(\beta_g)$ = 0.818 (0.74 to 0.90); 9,490 vehicle-days, 119 exits. Kinked: per extra $1,000 below market, $\exp(-\beta_-)$ = 1.30 (1.16 to 1.49); per $1,000 above market, $\exp(\beta_+)$ = 0.93 (0.76 to 1.11). Off-road trim against the reference trim: 3.72 (1.99 to 6.38). Per 10 points of discount off MSRP, in a companion model with discount in place of gap: 2.88 (1.79 to 4.41).

Worked (model). Under the single slope, a vehicle $2,145 above market has daily exit odds $0.818^{2.145}$ = 0.65 times those of an identical vehicle at market, which is about 54% longer expected time on the lot. Under the kinked model, the same comparison is $0.93^{2.145}$ = 0.86, with an interval that includes 1. The choice of model matters for any pricing recommendation, which is why the kinked estimate is reported.

A.7 Out-of-sample validation

Fit on vehicle-days with $t \le t_{\text{split}}$; predict $\hat{q}_{vt}$ on later days. Discrimination is the Mann–Whitney form of the AUC (Hanley and McNeil, 1982):

$$\text{AUC} = \left[\ \Sigma_{i:\,y=1}\ \text{rank}(\hat{q}_i) - n_1(n_1+1)/2\ \right] / (n_1\, n_0)$$

Calibration compares $\Sigma\,\hat{q}$ with $\Sigma\,y$ within thirds of g.

SPLIT	TRAIN VEHICLE-DAYS / EXITS	TEST VEHICLE-DAYS / EXITS	AUC WITH GAP	AUC WITHOUT GAP	PREDICTED / ACTUAL EXITS
Fit days 1–27, test 28–43	4,204 / 58	5,286 / 61	0.607	0.551	64.6 / 61
Fit days 1–31, test 32–43	5,173 / 73	4,317 / 46	0.581	0.559	58.5 / 46

> Robustness of the single-slope odds ratio (refit once each): drop dealer A 0.78; drop dealer B 0.84; drop dealer E 0.82; days before 37 only 0.78; without the month-end week 0.85; six-dealer days only 0.88; gap measured against the dealer's own lot median 0.76.

A.8 Revealed day cost and best-response price

Assume the daily exit hazard is approximately log-linear in price over the relevant range, $h(p) = h_0\, e^{-\beta p}$ with $\beta > 0$ per dollar, so expected days to exit are $T(p) = 1/h(p) = e^{\beta p}/h_0$. A dealer that values each day a vehicle sits at k dollars chooses p to maximize

$$\Pi(p) = p - k \cdot T(p) \qquad \Rightarrow \quad d\Pi/dp = 1 - k\beta T(p) = 0 \qquad \Rightarrow \quad k^* = 1 / (\beta\, T)$$

The second derivative, $-k\beta^2 T$, is negative, so this is a maximum. Rearranged for a dealer with a known k:

$$T^* = 1 / (\beta k) \qquad \text{cut } \Delta = \max(0, \ln(T / T^*) / \beta) \qquad \text{net gain} = k \cdot (T - \min(T, T^*)) - \Delta$$

> Plain English: at its best price, a dealer's price on a vehicle reveals how much it implicitly believes a day on the lot costs. If its true cost is higher, a cut pays; if lower, it does not. k should include floorplan interest, the risk of missing manufacturer volume targets, model-year ageing and the

> opportunity cost of the space. Assumptions: the hazard slope is known and constant; the cut does not change other margin (incentives, add-ons); competitors do not respond, so this is a unilateral best response, not a market equilibrium. If every dealer cuts together, relative price positions, and therefore the hazard gains, largely cancel.

Worked (synthetic). β = 0.0002 per dollar, T = 120 days. Implied k^* = 1/(0.0002 × 120) = $41.67 per day. With a true k = $60: T^* = 1/(0.0002 × 60) = 83.3 days; Δ = ln(120/83.3)/0.0002 = $1,823; net gain = 60 × (120 − 83.3) − 1,823 = $377 per vehicle. With k = $7.71 (a reported national average daily holding cost per new car, 2024): T^* = 648 days > 120, so Δ = 0. These inputs are invented for illustration and are not estimates for any dealer in the study.

A.9 Within-lot relative discount

$$\rho_{vt} = \tilde{\delta}_{vt} - \text{median}_{u \in d(v),\, t}\, \tilde{\delta}_{ut}$$

> Plain English: how much deeper (positive) or shallower a vehicle's smoothed discount is than the typical vehicle on the same lot that day. Because it compares a dealer only with itself, it is unaffected by how that dealer's website presents incentives. Exit rates by band use A.2 with U = vehicle-days.

Appendix B: Reproducibility protocol

1. **Define the competitive set.** Choose dealers and a model line. Record each dealer's public new-inventory URL.
2. **Collect daily.** At a fixed time each day, read every page of each dealer's results, follow pagination to the end, and store one dated file per dealer with VIN, advertised price, MSRP, trim, engine and drivetrain. Check each site's terms of use and robots.txt, and read at a rate that does not burden the site.
3. **Validate each snapshot.** Compare the number of vehicles read with the count the site displays. Mark the snapshot invalid if it read nothing or fewer than expected; never let an invalid snapshot imply exits.
4. **Build VIN histories** per Section 4.1, and classify exits with a guard rule of at least three snapshots.
5. **Smooth prices** with a three-snapshot median and detect dealer-wide moves (Appendix A.3).
6. **Run the synchronization test** on days when every dealer was observed (A.4).
7. **Fit the daily hedonic market price** and compute price gaps (A.5).
8. **Estimate exit rates and the exit model** with vehicle-level bootstrap intervals (A.2, A.6).
9. **Validate out of sample** on a time split, and refit with each dealer left out (A.7).
10. **Report** each claim with its evidence type, sample size and interval, using the wording rules in Section 7.

Data, confidentiality and methodology notes

Data. All inputs were public dealer web pages. The raw snapshots are not published, to protect the identities of the dealers studied. Aggregated tables in this paper are the complete basis for its claims.

Confidentiality. Dealers are labeled A to F in an order unrelated to their names. The market, brand, exact dates, VINs and dealer-level inventory counts are withheld so that dealers cannot be identified from combinations of details. Day numbers are counted from the first day of observation.

Software. Analyses were run in Python with pandas and NumPy. Logistic models use a ridge penalty of 1.0 on all coefficients except the intercept, fitted by Newton's method. Bootstrap intervals use 300 vehicle-level resamples; permutation tests use 20,000 shuffles.

Competing interests. ICONIC offers competitive-intelligence and AI services to automotive dealers, and the underlying data were collected in the course of that commercial work. No dealer reviewed or endorsed this paper.

Cite as. Christell, P. (2026). *Detecting Synchronized Repricing in Dealer Inventory: A VIN-Level Framework for Competitive Pricing Intelligence*. ICONIC Research. https://iconic.onl/research/synchronized-repricing-dealer-inventory/

References

1. Allison, P. D. (1982). Discrete-time methods for the analysis of event histories. *Sociological Methodology*, 13, 61–98.
2. AutoSuccess. (n.d.). LotLinx guides automotive industry on how to maximize used car inventory profits. autosuccessonline.com
3. Busse, M., Silva-Risso, J., & Zettelmeyer, F. (2006). $1,000 cash back: The pass-through of auto manufacturer promotions. *American Economic Review*, 96(4), 1253–1270.
4. Cachon, G. P., & Olivares, M. (2010). Drivers of finished-goods inventory in the U.S. automobile industry. *Management Science*, 56(1), 202–216.
5. Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.
6. Efron, B. (1979). Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7(1), 1–26.
7. Good, P. (2005). *Permutation, Parametric, and Bootstrap Tests of Hypotheses* (3rd ed.). Springer.
8. Hanley, J. A., & McNeil, B. J. (1982). The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1), 29–36.
9. Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34–55.
10. Trader. (n.d.). vAuto Provision: Market-based pricing. go.trader.ca
11. vAuto. (n.d.-a). Conquest: New car inventory management. vauto.com/products/conquest
12. vAuto. (2024). New car industry insights (daily holding cost per new car). vauto.com/industry-insights
13. White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4), 817–838.

> Busse et al. (2006) document how manufacturer promotions pass through to transaction prices, relevant to the incentive explanation in Section 5.1; Cachon and Olivares (2010) analyze dealer inventory levels, relevant to the days-of-supply measures.

© 2026 ICONIC. Research by Patrick Christell. This paper reports advertised listing data; it is not financial or legal advice.