

What Makes Dynamic Pricing Dynamic?

When and How to Update Prices Across a Network

Jörn Boehnke

Abstract

A retail chain must decide when to update prices and whether one rule can serve every location. We show how a shared rule can retain local timing without separate schedules. A fixed-time schedule permits increases at specified times. A gap trigger raises the price when the reference price exceeds the posted price by more than a threshold. Using 923 million timestamped fuel-price reports, we select all settings using 2024 data and test how closely the rules reproduce the 2025 diesel-price paths of 2,000 German stations, with a daily limit on increases and unrestricted cuts. Locally tuned schedules and triggers have nearly the same tracking error, but sharing their settings has very different effects. With one permitted increase per day, sharing adds 0.642 euro cents per liter to schedule error but only 0.097 to trigger error, a difference of 0.545. One shared trigger recovers 84 percent of the error reduction from replacing one shared schedule with 2,000 local schedules. Matching the shared trigger to within 0.10 cents per liter takes two schedule templates and station-specific assignments. The trigger’s sharing advantage persists within most brand networks. Two features explain the result: many stations can use the same threshold, and local price gaps determine when the trigger acts. With four permitted increases, a shared schedule that can wait for a positive gap outperforms 2,000 local schedules. Delayed information or limits that also count cuts can instead favor a fixed-time schedule over a trigger. A network can therefore share its update rule without imposing common update times.

Keywords: dynamic pricing; price updates; shared rules; state feedback; update timing; retail networks

1 Introduction

A retail chain needs to decide when to update prices and whether the same rule can serve every location. Pricing software may calculate a new price every minute, while the retailer changes the price customers see much less often. Updates can require approval, coordination, testing, and customer communication, and may also face software, contractual, or legal limits (Levy et al., 1997; Zbaracki et al., 2004). The retailer must decide which price changes to follow. Across a network,

it must also decide how much of that decision to customize. We study when and how to update prices when daily increases are limited.

The key distinction is between sharing a rule and sharing update times. We compare a fixed-time schedule with a gap trigger. Both take a reference price as given and decide when to update the posted price. A schedule permits increases at specified times of day. A trigger acts when the reference price exceeds the posted price by more than a threshold. Sharing a schedule gives every location the same possible increase times. Sharing a trigger gives every location the same threshold, but its own price gap determines when it acts. The rule is shared; the timing remains local. Using current local information to decide when to act is state feedback.

We test these rules using 923 million timestamped fuel-price reports covering 14,007 eligible German stations from 2019 through July 2026. For the main analysis, we reconstruct the 2025 diesel-price paths of 2,000 stations on a 15-minute grid. At this frequency, the panel captures not only price movements, but their timing across locations. This market combines frequent price changes, recurring cycles, and differences in timing across stations (Wang, 2009; Byrne and de Roos, 2019; Luco, 2019; Assad et al., 2024; Kahl and Wein, 2025). Each rule follows the same historical reference path at each station, subject to a daily limit on increases and unrestricted cuts. We select every setting using 2024 data and hold it fixed throughout the full 2025 evaluation year. Our outcome is tracking error: how closely each rule reproduces the observed 2025 price path.

Our main finding is that local accuracy does not tell a chain how well a rule will work when shared. Locally tuned schedules and triggers have nearly the same mean absolute tracking error. With one permitted increase per day, sharing adds 0.642 euro cents per liter to schedule error but only 0.097 to trigger error. The difference in sharing penalties is 0.545 cents per liter. One shared trigger recovers 84 percent of the improvement obtained by replacing one shared schedule with 2,000 local schedules. Current local information can therefore replace much of the need for local timing settings.

We next measure what it takes for schedules to close that gap. Two schedule templates come within 0.10 cents per liter of one shared trigger when each station is assigned a template using its own tuning history. A library of 32 schedules assigned using earlier station features does not. A shared trigger can therefore replace both a schedule library and the station-specific assignments needed to use it. The trigger’s sharing advantage persists within most brand networks, where the average difference in sharing penalties is about half the national value.

Many stations can use the same threshold with little loss of accuracy, while their current price gaps determine local update times. A simple model shows why one threshold can work across locations even when they update at different times. The daily limit also matters. A waiting schedule keeps an increase opportunity available until a positive price gap occurs. With one permitted increase, waiting recovers about half the error reduction from replacing a shared schedule with a shared trigger. With four, the waiting schedule has the lowest error among the shared rules tested and outperforms 2,000 local schedules by 0.075 cents per liter.

This question connects several literatures. Dynamic-pricing research studies which prices firms

should choose (Gallego and van Ryzin, 1994; Fisher et al., 2018), and algorithmic-pricing research examines how software changes those choices (Miklós-Thal and Tucker, 2019; Harrington, 2022; den Boer et al., 2026). Work on limited adjustments and fixed calendars studies how much can be achieved with a few well-timed changes (Chen et al., 2016; Ma et al., 2021). Our schedule-trigger comparison also relates to periodic and continuous inventory review, time- and state-dependent price adjustment, and event-triggered control (Zipkin, 2000; Caplin and Spulber, 1987; Klenow and Kryvtsov, 2008; Heemels et al., 2012). We study how accuracy changes when a network replaces local settings with shared ones. This connects update timing with research on transferring organizational practices across locations (Darr et al., 1995; Jensen and Szulanski, 2007). Related work examines how information and automation enable consumers to respond to dynamic prices (Bollinger and Hartmann, 2020).

With current information and unrestricted cuts, a shared rule can preserve local responsiveness without separate local schedules. Delayed information or limits that also count cuts can instead favor schedules over triggers. A network can therefore standardize how locations decide when to update without requiring them to update at the same times.

2 Price updates across a network

2.1 Retailer choice under a daily increase limit

Consider a retailer with many locations, or a vendor that manages prices on its behalf. A pricing engine, a cost-based rule, or a manager supplies a reference price for each location. The price-update rule decides when to move the posted price to the current reference price. We take the reference prices as given and study when they are implemented.

In the main analysis, each station may raise its price at most K times per day. Cuts are unrestricted and occur whenever the reference price falls below the posted price. Whenever a rule increases the price, it posts the current reference price. It cannot post above the reference price in anticipation of later increases. The limit applies separately at each station rather than to a pool of increases shared across the network.

Germany provides a concrete example of a daily increase limit. Fuel stations report price changes to the Market Transparency Unit for Fuels (Bundeskartellamt, 2026). The 2026 law permits one increase at noon and leaves cuts unrestricted. It applies to station operators and companies that prescribe their prices, so a central price setter must comply at each station (Kraftstoffpreisanpassungsgesetz, 2026). Austria has restricted increases in the same one-sided way since 2011, while Western Australia has required posted prices to remain in effect for a full day since 2001 (Obradovits, 2014; Wang, 2009).

Recent empirical work studies how Germany’s rule affects prices and margins (Jung et al., 2026; Frondel et al., 2026). Related theory examines how one-sided adjustment limits affect pricing and welfare (Denter, 2026; Paarsch and Schmedders, 2026) and where a once-daily reset should occur (Siemroth, 2026). Our question is different: given a reference-price path and a daily limit

on increases, when should those increases occur, and how much of that timing rule can be shared across locations? Germany’s law motivates the one-sided limit. We vary the timing to study similar technical or organizational limits in which firms can choose when to update.

2.2 Schedules, triggers, and local tuning

A fixed-time schedule specifies K times of day. At a scheduled time, it raises the posted price to the reference price if the reference price is higher than the posted price. Otherwise, that opportunity expires. A gap trigger specifies one threshold. At each quarter-hour, it calculates the price gap: the reference price minus the posted price. It raises the price to the reference when the gap strictly exceeds the threshold and the daily increase limit has not been reached.

Both rules observe the same reference path and keep track of their own posted price. A schedule fixes possible increase times in advance, while a trigger uses its current price gap to decide when to act. The trigger therefore needs no additional price variable.

A waiting schedule also specifies K times. Each time releases one increase opportunity, which remains available until it is used or the day ends. The rule uses an available opportunity at the first eligible quarter-hour with a positive price gap. Unlike a fixed-time schedule, it does not immediately discard an unused opportunity. Unlike a trigger, it retains scheduled releases and does not use gap size to decide when to act. A hybrid reserves at most two of the K opportunities for scheduled times and assigns the rest to a trigger. Appendix B.4 describes these rules in detail.

Each rule can have local or shared settings. Local settings are selected separately for each station. A shared setting is selected once for all stations in the comparison. Shared settings do not impose a common price, and a shared trigger does not impose common update times. We measure the additional tracking error from replacing local settings with shared settings. We call this the sharing penalty.

Let $L_{f,c,K}$ denote tracking error for rule f , calibration level c , and daily increase limit K . The rules are schedule S and trigger T ; calibration is common C or local L . Their sharing penalties are

$$P_{S,K} = L_{S,C,K} - L_{S,L,K}$$

and

$$P_{T,K} = L_{T,C,K} - L_{T,L,K}.$$

The difference in sharing penalties is

$$\Delta_K = P_{S,K} - P_{T,K}.$$

A positive Δ_K means that sharing adds more error to schedules than to triggers. All four errors are needed. Comparing only the two shared rules would mix differences already present under local settings with differences introduced by sharing. We therefore report the shared-rule comparison alongside the difference in sharing penalties.

2.3 Why shared rules can retain local timing

A two-opportunity example shows the trade-off. Each station starts at a baseline price with one increase available. In an early quarter-hour, the reference price rises by $g > 0$ and then returns to baseline, so unrestricted cuts restore the posted price. In a later quarter-hour, the reference price rises by $H > g$ with probability $\rho \in (0, 1)$; otherwise it stays at baseline. The later event is unknown when the first decision is made.

Using the increase early leaves expected later tracking error ρH . Waiting leaves current error g but preserves the increase for later. The station therefore acts early when

$$g > \rho H.$$

A gap threshold $h = \rho H$ implements this decision. It captures the opportunity cost of using a scarce increase now rather than preserving it for a potentially larger gap later.

Now consider two stations with the same g , H , and ρ but different episode times, with $0 < g < \rho H$ and all four episode times distinct. Each local schedule places its increase at its own later episode. Any threshold $h \in [g, H)$ makes the same choices. Local schedules and local triggers therefore have the same expected tracking error. The same threshold also works at both stations, so sharing does not increase trigger error. A shared schedule has only one clock time and can cover the later episode of at most one station. Equal local accuracy can therefore coexist with different sharing penalties.

With two increases, both episodes can be corrected in this example. A waiting schedule can do the same when both opportunities are released before the first episode. With one increase, a waiting schedule released before the early episode acts on the first positive gap, while a threshold can pass up that gap and preserve the increase for later. The example separates two decisions: preserving an unused opportunity and deciding whether the current gap is large enough to use it. We test both in the data.

3 Data and replay design

3.1 Price records and sample

We use 923,031,338 Tankerkönig price reports from January 2019 through July 2026 (Tankerkoenig, 2026). Each report records a timestamp, station identifier, and posted prices for diesel, E5, and E10. We reconstruct the price in effect at each station on a 15-minute grid in Europe/Berlin time. The main evaluation contains 2,000 diesel stations and 725,562 station-days in 2025.

We divide eligible stations into four groups by their earlier rate of price changes and sample 500 from each group. This ensures that both less active and more active stations are represented. The largest standardized difference between the sample and the 14,007-station eligible frame is 0.062. We retain station-days with little or no price activity. Appendix A.2 gives the eligibility rules and sample comparisons.

3.2 Reference path and tracking error

Let p_{it}^0 be the historical posted price at station i in quarter-hour t . We use this observed sequence as the reference path. Let q_{ait} be the price posted by rule a in the replay. Mean absolute tracking error is

$$L_a = \frac{1}{N_a} \sum_{(i,t) \in \mathcal{E}_a} |q_{ait} - p_{it}^0|.$$

Here, $\mathcal{E}_a$ contains the eligible station-quarter-hour observations and N_a is their number. Rules being compared use the same observations. Error is measured in euro cents per liter; lower values mean closer tracking.

The reference paths are observed posted prices, not firms’ internal targets, and may already reflect existing posting frictions and competitors’ responses. We hold them fixed and ask how closely each rule reproduces them under a daily increase limit. Tracking error therefore measures execution, not profit, welfare, or the causal effect of the German law. For a retailer applying the same analysis to its own target-price stream, squared deviations from the target locally approximate losses in a smooth business objective, with curvature determining their weight.¹ Section 4.5 reports the same ranking under root mean squared error. With current information and unrestricted cuts, the replayed price never exceeds the reference, so every error is a shortfall. That property need not hold when information is delayed or cuts are restricted. We also report the share of observations within one cent of the reference and measures of large tracking errors.

The main measure gives equal weight to each quarter-hour. We also use traffic weights based on a national weekday-by-hour profile from official counters on Autobahns and Bundesstraßen (Bundesanstalt für Straßen- und Verkehrswesen (BASt), 2026). Every station receives the same 168-cell profile, with weights normalized within each station-day. This changes which times matter most, not which station has the greatest traffic. The counts measure passing vehicles, not station visits or fuel purchases. We keep the original rule settings and change only the evaluation weights.

3.3 Selection, evaluation, and inference

We use 2019–2023 data to describe earlier station behavior and 2024 data to select settings. Shared settings are selected using the pooled tuning sample; local settings are selected separately by station. Five stations without eligible 2024 tuning observations use fallback local settings (Appendix A.2). We keep all settings fixed during the 2025 evaluation.

The main comparisons allow at most one or four increases per day. One creates the tightest timing problem; four allow more corrections. The appendix also reports limits of two and eight. The posted price carries across midnight, while the daily increase allowance resets. After selecting the settings, we replay each rule through 2024 to initialize its posted price at the start of 2025. Appendix B.9 gives the action order and time conventions.

¹If r_{it} is such a target and locally maximizes a twice-differentiable objective $\pi_{it}(p)$, then $\pi_{it}(r_{it}) - \pi_{it}(q_{it}) = -\frac{1}{2}\pi_{it}''(r_{it})(q_{it} - r_{it})^2 + o((q_{it} - r_{it})^2)$.

For a schedule with one permitted increase, we evaluate all 96 quarter-hour slots. At higher limits, the main analysis adds one slot at a time by forward selection. As a check, we also select the shared four-increase schedule exactly by dynamic programming using 2024 data. That schedule lowers 2025 tracking error by 0.064 cents per liter relative to forward selection. We also select shared and local schedules exactly on a fixed 200-station subsample. These checks assess whether the schedule-search method drives the results.

Trigger thresholds range from 0.25 to 8.00 cents per liter in steps of 0.25. Nearby thresholds often produce the same actions because recorded price gaps often move in whole-cent increments. Local one-increase triggers select 17 of the 32 candidates.

Stations can face the same national price movements on a given date. We use 1,999 bootstrap draws that resample stations within the four activity groups and resample dates using seven-day moving blocks. All rules within a comparison use the same draws. Confidence intervals are simultaneous within the stated comparison families. They condition on the settings selected using 2024 data; the selection step is not resampled. Appendix B.8 gives the inference procedure.

3.4 Tests of the mechanism

We define a station’s repricing pattern as the distribution of its upward price movements across weekdays and quarter-hours. Using 2024 data, we sum the positive price changes in each weekday-quarter-hour cell and normalize the 672 values to sum to one. The network pattern is the equally weighted average of station patterns. Pattern distance is half the sum of absolute differences between a station’s pattern and that average, the total-variation distance between the two distributions.

We call a setting acceptable when its tracking error is within the stated tolerance of that station’s lowest candidate error. For the one-increase rules, we evaluate every station under all 96 schedule times and 32 thresholds in both 2024 and 2025. We compare error at the shared 2024 setting with each station’s best candidate in the relevant year. The 2025 comparison is descriptive; all rule settings still come from 2024.

Two tests change the information used to decide increases. One uses the station’s reference path from exactly 364 days earlier, adjusted to its current opening price. The other circularly shifts, within each day, the gap signal generated by a rule that follows cuts but makes no increases. In both tests, cuts and executed increases still use the current reference price. Only the decision signal changes. Each altered signal is compared with a current-signal rule on the same observations. Appendix G.2 gives the details. Because a circular shift can move later information earlier, we use it only as a timing test.

Finally, we change the reference paths themselves. One transformation smooths short-run movements using a centered three-hour rolling median. Another subtracts each station’s mean quarter-hour-of-day price profile estimated from 2019–2023. Removing this recurring intraday level pattern is not the same as removing the station’s full weekly repricing pattern. For each transformed path, we compare settings retuned on transformed 2024 data with settings retained from the original

path. We also calculate the lowest feasible within-day tracking error attainable with knowledge of the remaining day’s prices, starting from the same posted price as the evaluated rule. We call the difference between a rule’s error and this benchmark its hindsight gap. Appendix G.4–G.6 gives the timing tests, transformations, and benchmark.

4 Results

4.1 One shared trigger nearly matches local schedules

Local schedules and triggers have nearly identical tracking error: 2.079 and 2.086 cents per liter, respectively, with one permitted increase per day. Sharing their settings changes the picture. Schedule error rises to 2.721, while trigger error rises only to 2.183 (Table 1). Figure 1 shows the same pattern as settings move from local to brand-network to one rule shared across the full sample.




Figure 1: Tracking error as local rules are replaced by a brand-network rule and by one shared rule. Panels show one and four permitted daily increases under equal time weights. Schedule lines use squares; trigger lines use circles. Rules are selected using 2024 data and evaluated over full-year 2025 on 2,000 stations. The brand-network point tunes one rule within each of 16 recorded brand networks and pools the remaining stations.

Sharing therefore adds 0.642 cents per liter to schedule error and only 0.097 to trigger error. The difference in sharing penalties is 0.545 cents per liter. Rules that are nearly tied when tuned locally therefore respond very differently to sharing.

For a network, the gain is large. Replacing one shared schedule with one shared trigger captures 84 percent of the error reduction obtained by replacing that shared schedule with 2,000 local schedules. The shared trigger remains only 0.104 cents per liter above those local schedules.

The result persists with four permitted increases. The difference in sharing penalties remains 0.332 cents per liter, and the shared trigger is only 0.052 cents per liter above local schedules. Meanwhile, local schedule-minus-trigger differences remain near zero at both limits: -0.007 cents at one increase and 0.017 at four. Section 4.3 asks whether allowing a schedule to wait can improve

Table 1: Full-year tracking error and sharing penalty

K	Quantity	Schedule	Trigger	Schedule – Trigger
1	Shared-rule error	2.721	2.183	0.538 [0.479, 0.597]
1	Local-rule error	2.079	2.086	-0.007 [-0.037, 0.022]
1	Sharing penalty	0.642	0.097	0.545 [0.486, 0.604]
4	Shared-rule error	1.134	0.785	0.349 [0.322, 0.375]
4	Local-rule error	0.734	0.717	0.017 [-0.006, 0.039]
4	Sharing penalty	0.400	0.068	0.332 [0.300, 0.364]

Notes: Entries are cents per liter over full-year 2025. Rules are selected using 2024 data and evaluated on 2,000 stations. The sharing penalty is shared minus local error within a rule form. The final column reports schedule minus trigger for each quantity. Brackets are simultaneous 95 percent intervals from bootstrap draws that resample stations within activity groups and dates using seven-day moving blocks, with the 2024 settings held fixed.

further.

The main contrast remains large on all 14,007 eligible stations, in three earlier evaluation years, under root mean squared error, and with federal-road traffic weights. Section 4.5 reports these and the remaining sample, period, and measurement checks.

4.2 How much local configuration a single shared trigger replaces

We next ask how much local configuration a schedule-based system needs to match one shared trigger. Figure 2 compares schedule libraries that vary in the number of templates and in how stations are assigned to them. A template is a set of update times that can be assigned to multiple stations.




Figure 2: How much customization one shared trigger replaces. Panel A: full-year tracking error for libraries of G schedule templates or thresholds, with one permitted increase. Filled marks assign each station its own 2024-best template from a set chosen on the stations’ tuning losses; open marks assign templates from pre-period station features. The horizontal line is one shared trigger. Panel B: the same with four permitted increases; schedule templates are drawn from 200 locally tuned four-time schedules. Panel C: the distribution across 2,000 stations of shared-trigger minus local-schedule error with one permitted increase.

First, we use each station’s own 2024 price history to build the library and assign its best template. With one permitted increase, two schedule templates come within 0.10 cents per liter of one shared trigger. Four slightly outperform it in point estimates; with eight, the advantage is statistically clear. With four permitted increases, four templates reach the 0.10-cent tolerance, sixteen slightly outperform the trigger in point estimates, and 32 are the first to do so with an interval entirely below zero. Appendix Tables A11–A12 report the corresponding intervals.

Second, we group stations using 2019–2023 features and select settings on 2024 data within each group. No library of up to 32 schedule templates comes within 0.10 cents per liter of the shared trigger at either daily limit. Earlier station features do not match what station-specific tuning achieves. At four increases, the own-history library uses templates from the 200-station exact-schedule sample, so the two curves differ in how templates are built as well as assigned. Appendix E.8 gives the details.

One shared trigger can therefore replace both a library of schedules and the station-specific assignments needed to use it. It replaces not only many clock times, but also the need to decide which schedule each station should use.

The sharing advantage is also widespread across stations. With one permitted increase, the per-station difference in sharing penalties is positive at 80 percent of stations. The shared trigger has lower error than the local schedule itself at 38 percent of stations and is within 0.5 cents per liter of it at 73 percent. It trails the local schedule by more than 0.5 cents at 18 percent of stations and by more than one cent at only 1.5 percent. The sharing-penalty difference is positive in all four activity groups, ranging from 0.923 cents in the least active group to 0.169 in the most active. The Online Appendix reports the four cells by activity group.

At one permitted increase, the trigger’s advantage comes mainly from shrinking large errors rather than increasing the fraction of time close to the reference. The shared schedule and shared trigger are each within one cent of the reference for about 55 percent of evaluated quarter-hours, but the trigger reduces the 95th-percentile error from 11.0 to 8.0 cents per liter. With four increases, it improves both measures: the within-one-cent share rises from 77 to 82 percent, while the 95th-percentile error falls from 5.8 to 5.0 cents. Appendix Table A14 reports the full error-size and duration comparisons.

The sharing advantage also persists within recorded brand networks. Across the 16 brand networks with at least 20 sample stations, the pooled difference in sharing penalties is 0.276 cents per liter at one increase and 0.169 at four. The difference is positive in 14 of 16 brand networks at one increase, and the brand-network trigger ends only 0.073 cents above local schedules. Across brand networks, the sharing advantage rises with dispersion in member stations’ repricing patterns (rank correlation 0.45). Omitting any one brand changes the national contrast by at most 0.072 cents. The brand-network results therefore show that the sharing advantage persists within narrower, brand-defined networks. Appendix E.10 also reports results by recorded parent group.

4.3 How the best shared rule changes with the daily limit

A trigger is not the only way to preserve local timing. Table 2 adds the waiting schedule and hybrid to the comparison. The waiting schedule keeps a scheduled increase opportunity available until a positive price gap occurs. It therefore shows how much is gained by preserving an unused opportunity rather than letting it expire immediately.

Table 2: Shared and local tracking error for four rule forms

K	Rule	Shared	Local	Sharing penalty	Unused (%)
1	Fixed-time schedule	2.721	2.079	0.642 [0.575, 0.709]	17
1	Waiting schedule	2.464	2.049	0.414 [0.370, 0.458]	0.7
1	Gap trigger	2.183	2.086	0.097 [0.077, 0.116]	-
4	Fixed-time schedule	1.134	0.734	0.400 [0.364, 0.437]	18
4	Waiting schedule	0.659	0.568	0.091 [0.076, 0.106]	10
4	Gap trigger	0.785	0.717	0.068 [0.057, 0.080]	-
4	Hybrid	0.698	0.597	0.101 [0.087, 0.115]	15

Notes: Entries are cents per liter over full-year 2025 on 2,000 stations. The waiting schedule releases an increase opportunity at a stored time and uses it at the next positive gap before the end of the day. The hybrid reserves at most two opportunities for scheduled times and leaves the rest to one threshold. Every rule is selected using 2024 data. Unused is the share of scheduled increase opportunities that expire unused. The shared trigger makes no increase on 3.2 percent of station-days when one increase is permitted. Brackets on the penalties are simultaneous 95 percent intervals within each daily limit.

With one permitted increase per day, waiting helps substantially but does not match the trigger. The shared fixed-time schedule makes 0.83 increases per day on average, compared with 0.99 for the waiting schedule. Waiting recovers 48 percent of the error reduction from replacing the shared schedule with the shared trigger. Its sharing penalty is still 0.414 cents per liter, 0.318 higher than the trigger’s. The selected shared release time is midnight, so the waiting schedule acts at the first positive gap of the day. Under this tight limit, using gap size to decide when to act is more accurate than acting on the first eligible positive gap.

With four permitted increases, the waiting schedule becomes the most accurate shared rule we test. Its error is 0.659 cents per liter, below the trigger’s 0.785 and the hybrid’s 0.698. It also outperforms 2,000 local fixed-time schedules by 0.075 cents per liter. Its sharing penalty exceeds the trigger’s by only 0.022 cents per liter. At this limit, preserving an opportunity until the gap turns positive is enough to outperform rules that condition on gap size.

The additional limits sharpen the pattern. At two increases, the hybrid has the lowest shared-rule error. At eight, the waiting schedule outperforms the trigger by 0.030 cents per liter. Their sharing penalties differ by only 0.002 cents per liter. Appendix G.8 reports all four limits.

Together, these results show how the value of price-gap information changes with the daily limit. With one increase, gap size is valuable because using an opportunity on a small gap can leave no opportunity for a larger one later. With more increases, preserving an opportunity until any positive gap appears becomes increasingly effective. The selected shared threshold moves in the same direction, falling from 4.25 to 3.25, 1.50, and 1.25 cents per liter as the limit rises from one to two, four, and eight increases. More opportunities let the trigger act on progressively smaller

gaps.

4.4 Why the shared trigger works

A shared trigger needs a threshold that works across stations and local information that tells it when to act. Figure 3 separates these two requirements and shows why schedules are harder to share.




Figure 3: Why the shared trigger works. Panel A: share of stations for which the shared 2024 setting is within a stated error tolerance of the station's best 2025 setting, for thresholds and schedule times with one permitted increase. Panel B: shared-rule error and sharing penalty for triggers using the current, 364-day-lagged, or circularly shifted gap, evaluated on the same observations. Panel C: schedule and trigger sharing penalties by quartile of repricing-pattern distance at four permitted increases. Panel D: four-increase difference in sharing penalties on the original, smoothed, and de-cycled paths after retuning, and when the original settings are held fixed. Brackets are simultaneous 95 percent intervals within each analysis.

Thresholds overlap much more across stations than schedule times. At one permitted increase, the shared threshold produces error within 0.10 cents per liter of a station's best 2025 candidate at 63 percent of stations, compared with 14 percent for the shared schedule time. At a 0.25-cent tolerance, the shares are 77 and 30 percent (Figure 3A). The key is therefore not that stations share

the same best threshold, but that the same threshold remains accurate across many locations. The trigger retains a sharing advantage when schedules and triggers are restricted to the same number of candidate settings (Appendix G.1).

The live price gap then determines whether that shared threshold acts at useful local times. The current and 364-day-lagged triggers have similar sharing penalties, 0.098 and 0.092 cents per liter, yet the current shared rule has 0.158 cents less error when both are evaluated on the same observations. A circularly shifted signal has an even smaller sharing penalty, 0.008, but produces 0.370 cents more shared-rule error than the same trigger using the current gap (Figure 3B). Overlapping settings therefore explain why the trigger is easy to share; the live gap explains why it remains accurate once shared.

Schedule sharing becomes costly when stations follow different repricing patterns. With four permitted increases, the schedule penalty rises from 0.220 cents per liter among the quarter of stations closest to the network pattern to 0.714 among the farthest quarter (Figure 3C). Trigger penalties move much less, from 0.070 to 0.099. The sharing-penalty difference rises by 1.336 cents per liter per unit of repricing-pattern distance, while the trigger’s own slope is only 0.047. A constructed-path test shows the same timing effect: spreading fitted station patterns farther apart increases the four-increase sharing-penalty difference. Appendix G.3–G.4 reports the full timing tests.

Removing recurring features of the price path does not remove the result. After subtracting each station’s recurring daily price profile and retuning every rule, the four-increase sharing-penalty difference remains 0.284 cents per liter, compared with 0.333 on the original paths for the same 1,995 stations (Figure 3D). At one increase, it remains 0.326, compared with 0.546 on the original paths. Smoothing the path likewise leaves a four-increase difference of 0.222. Holding the original settings fixed on the de-cycled path instead produces -0.053. Retuning therefore leaves most of the four-increase sharing advantage and more than half of the one-increase advantage. The sign reverses only when the original settings are carried over unchanged. Appendix G.5 reports the full transformed-path comparisons.

4.5 Sharing advantage across benchmarks, years, and samples

Table 3 shows the sharing-penalty difference under alternative outcomes, benchmarks, years, and samples.

The same pattern appears in the full station frame and in earlier years. Using all 14,007 eligible stations gives a one-increase difference in sharing penalties of 0.555 cents per liter, almost identical to the main estimate of 0.545. In three earlier evaluation years, with settings selected on the previous calendar year, the one-increase estimates range from 0.381 to 0.443. These replications evaluate April through July rather than a full calendar year.

Traffic weights and root mean squared error give the same ranking. Federal-road traffic weights increase the one-increase difference in sharing penalties to 0.890 cents per liter and leave a four-increase difference of 0.208. Under root mean squared error, which gives more weight to large

Table 3: Difference in sharing penalties across analyses

Analysis	Stations	$K = 1$	$K = 4$
Full-year tracking	2,000	0.545 [0.486, 0.604]	0.332 [0.300, 0.364]
Full eligible frame	14,007	0.555 [0.519, 0.591]	0.328 [0.311, 0.345]
Federal-road exposure	2,000	0.890 [0.816, 0.965]	0.208 [0.180, 0.237]
Hindsight gap	2,000	0.780 [0.720, 0.840]	0.369 [0.337, 0.401]
Projected reference	200	0.974 [0.691, 1.256]	0.475 [0.355, 0.595]
Exact schedules	200	0.624 [0.449, 0.798]	0.278 [0.173, 0.384]
2022, tuned on 2021	1,923	0.399 [0.315, 0.483]	0.096 [0.069, 0.123]
2023, tuned on 2022	1,954	0.443 [0.371, 0.515]	0.352 [0.327, 0.377]
2024, tuned on 2023	1,986	0.381 [0.304, 0.458]	0.234 [0.207, 0.261]
Opening hours, spring	1,162	1.136 [0.951, 1.321]	0.186 [0.141, 0.231]
E5 grade, spring	1,964	0.469 [0.398, 0.540]	0.270 [0.244, 0.296]
E10 grade, spring	1,902	0.445 [0.372, 0.518]	0.249 [0.223, 0.275]
Root mean squared error	2,000	0.931 [0.853, 1.008]	0.638 [0.571, 0.706]

Notes: Entries are the schedule sharing penalty minus the trigger sharing penalty, in cents per liter. Positive values mean that sharing penalizes schedules more than triggers. Full-year tracking repeats Table 1. The full eligible frame uses all eligible stations. Federal-road exposure weights quarter-hours by the vehicle-passage profile. The hindsight gap is each rule’s error above the lowest within-day error attainable from the same starting state with knowledge of the remaining reference path. The projected reference uses a path that obeys the daily limit on a fixed 200-station sample and evaluates the original settings on that path. If the rules are reselected on the projected path, the shared trigger reproduces it exactly; the appendix reports that case. Exact schedules replace forward-selected schedules on the same sample. Earlier-year rows select settings on the prior year and evaluate April through July. Opening-hours and fuel-grade rows use the same spring evaluation window. Brackets are simultaneous 95 percent intervals within each analysis.

deviations, the differences are 0.931 and 0.638. Under this metric, local triggers also outperform local schedules at one increase, so the trigger’s advantage extends beyond the sharing margin.

Exact schedule search and the two alternative benchmarks give the same ranking. Selecting shared and local schedules exactly on the fixed 200-station sample leaves a four-increase difference in sharing penalties of 0.278 cents per liter. Relative to each rule’s own within-day hindsight benchmark, the differences are 0.780 at one increase and 0.369 at four. On a constructed reference path that itself satisfies the daily increase limit, they are 0.974 and 0.475.

The same ranking holds during recorded opening hours and for E5 and E10. At one increase, the differences are 1.136, 0.469, and 0.445 cents per liter, respectively. All three use the same spring evaluation window. The shared trigger also has lower error than the shared schedule on stations not used to fit either shared rule. Appendix E.3 reports the held-out-station results.

5 Choosing a rule for the network

5.1 When a shared fixed-time schedule beats the trigger

The ranking of the shared fixed-time schedule and shared trigger depends on the information available to the rule and on which price changes count toward the limit. Table 4 changes each in turn.

A delay directly weakens the trigger’s ability to use current local information. With a one-hour

Table 4: When a shared fixed-time schedule can outperform a shared trigger

Condition	K	Schedule – Trigger	
		Shared-rule error	Sharing penalty
Current information; increases only	1	0.538 [0.479, 0.597]	0.545 [0.486, 0.604]
Current information; increases only	4	0.349 [0.322, 0.375]	0.332 [0.300, 0.364]
15-minute delay; increases only	8	0.390 [0.372, 0.407]	0.255 [0.234, 0.276]
60-minute delay; increases only	8	-0.161 [-0.179, -0.144]	0.111 [0.094, 0.127]
Current information; all changes count	1	-0.495 [-0.545, -0.445]	0.147 [0.122, 0.173]
Current information; all changes count	4	-0.142 [-0.185, -0.098]	0.037 [-0.005, 0.080]

Notes: Entries are cents per liter over full-year 2025 on 2,000 stations. Both numerical columns report schedule minus trigger. Negative shared-error values favor the schedule; positive sharing-penalty values mean that sharing penalizes schedules more. Rules are reselected using 2024 data for each condition. These comparisons rerun the fixed-time schedule and gap trigger; waiting schedules and hybrids are not reselected. Brackets are simultaneous 95 percent intervals.

delay and eight permitted increases, the shared schedule has 0.161 cents per liter less error than the shared trigger. The delayed eight-increase trigger has about the same error as a one-increase trigger with current information: 2.24 versus 2.18 cents. In this comparison, a one-hour delay wipes out the gain from seven additional permitted increases. The timing rule and the timeliness of its price information must therefore be considered together.

Counting cuts can also reverse the ranking. Under this limit, the trigger acts in either direction: it posts the current reference price when the absolute price gap exceeds its threshold and a daily change opportunity remains. It may therefore use its only change on an early cut and be unable to respond to a later increase. The shared schedule has 0.495 cents per liter less error at one permitted change and 0.142 cents less at four. Appendix H.2 gives the rules used under this limit.

Error levels and sharing penalties answer different questions. A schedule can have lower shared-rule error even when sharing penalizes it more than the trigger. At four permitted changes under the all-change limit, the difference in sharing penalties is only 0.037 cents per liter, so the two sharing penalties are nearly tied.

Fixed-time schedules also provide a benefit outside tracking accuracy: their possible increase times are known in advance. A trigger’s increase times depend on the realized price gap. Predictable timing may matter when retailers must coordinate or communicate changes, and advance price commitments can affect consumer responses (Kabul and Parlaktürk, 2019).

5.2 Operational implications

For a network that wants accurate tracking with little local configuration, the central choice is whether current local price gaps can determine update times. With current information and unrestricted cuts, the most accurate shared rule changes with the daily limit: the gap trigger at one increase, the hybrid at two, and the waiting schedule at four and eight. When information is delayed or cuts also consume capacity, a fixed-time schedule can instead be more accurate. The rule, the information feed, and the change limit should therefore be chosen together.

Local tuning is most valuable where common settings fit poorly. Matching one shared trig-

ger with schedules requires both a library of schedule templates and station-specific assignments. Repricing-pattern distance is the strongest measured correlate of the schedule sharing penalty: locations with more unusual timing patterns lose more from sharing common schedule times. The configuration problem is therefore not simply how many local schedules to maintain, but where local timing is valuable enough to warrant them.

Settings also differ in how stable they are over time. Only 55 percent of stations keep the same one-increase schedule time from the 2023 to 2024 tuning year. On 2025 data, retaining a station’s 2024 setting rather than choosing its best 2025 candidate adds 0.150 cents per liter for schedules and 0.063 for triggers. The shared fixed-time schedule also lets 17 percent of its one-increase opportunities expire unused, while the shared trigger makes no increase on only 3.2 percent of station-days. Appendix I.3 reports the full stability and update-use results.

A retailer can run the same analysis on its own reference-price stream and operating constraints. With its own sales volumes, margins, and demand responses, it could evaluate the rules against its own business objective. Appendix F.6 uses German diesel sales to show the scale under a constant-volume assumption.

5.3 Update limits can change the reference path

Changing the update limit can change the prices firms choose, not just when they post them. The observed 2026 adjustment pattern differs sharply from a mechanical replay of the 2025 paths under the noon rule. On the common 15-minute grid, the observed 2026 paths average 8.07 cuts per station-day from April 1 through June 17, compared with 1.91 in the corresponding replay. Appendix J reports the raw-event counts, alternative time resolutions, and separate periods. This comparison cannot separate changes in target prices, starting prices, costs, competitors’ responses, or other market conditions.

Theory gives one concrete mechanism. When another increase will be unavailable, a firm may post above its current target to leave room for later cuts (Paarsch and Schmedders, 2026; Denter, 2026; Siemroth, 2026). Our replay cannot use that response because it never posts above the reference price. The main analysis instead holds the reference path fixed and asks which rule follows it most closely under the change limit. In practice, the update rule and the pricing process that generates the reference path should be evaluated together.

6 Conclusion

A network does not need a separate schedule at every location to update prices at different local times. In our German fuel-price replays, local schedules and triggers have nearly identical mean absolute tracking error, but sharing their settings has very different effects. With one permitted increase per day, a shared trigger recovers 84 percent of the error reduction from replacing one shared schedule with 2,000 local schedules. Matching the shared trigger to within 0.10 cents per liter takes two schedule templates and station-specific assignments. The trigger’s sharing advantage

also persists within most brand networks.

The mechanism is simple. Many stations can use the same threshold with little loss of accuracy, while each station’s current price gap determines when that threshold is crossed. When increases are scarce, gap size helps determine whether an opportunity is worth using. As the limit rises, preserving an opportunity until a positive gap appears becomes more effective: with four permitted increases, the shared waiting schedule is the most accurate shared rule we test and outperforms 2,000 local fixed-time schedules. Delayed information or limits that also count cuts can instead make a fixed-time schedule more accurate.

The same problem arises whenever a firm computes recommended prices more often than it can implement them. Our replay holds the reference path fixed and compares how update rules follow it; changing the update constraint can also change the prices a firm chooses to generate. The update rule and the pricing process that supplies the reference path should therefore be designed together. Standardizing a pricing rule need not standardize the timing of price changes: a network can share how locations decide when to act while letting local conditions determine when they act.

Data and Code Availability

The replication files include analysis code, nonconfidential aggregate data, software setup files, and instructions for obtaining the source data. Raw Tankerkönig records are not redistributed and are available from the original source (Tankerkoenig, 2026); BAST traffic data are publicly available (Bundesanstalt für Straßen- und Verkehrswesen (BAST), 2026).

References

- Assad, S., Clark, R., Ershov, D., and Xu, L. (2024). Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market. *Journal of Political Economy*, 132(3):723–771.
- Bollinger, B. K. and Hartmann, W. R. (2020). Information vs. automation and implications for dynamic pricing. *Management Science*, 66(1):290–314.
- Bundesanstalt für Straßen- und Verkehrswesen (BAST) (2026). Automatische Zählstellen auf Autobahnen und Bundesstraßen: Stundenwerte. Official data portal and annual hourly archives. Final 2023 and 2024 archives; CC BY 4.0; accessed August 31, 2026.
- Bundeskartellamt (2026). Market transparency unit for fuels. Official information page. Accessed August 12, 2026.
- Byrne, D. P. and de Roos, N. (2019). Learning to coordinate: A study in retail gasoline. *American Economic Review*, 109(2):591–619.
- Caplin, A. S. and Spulber, D. F. (1987). Menu costs and the neutrality of money. *Quarterly Journal of Economics*, 102(4):703–725.

- Chen, Q. G., Jasin, S., and Duenyas, I. (2016). Real-time dynamic pricing with minimal and flexible price adjustment. *Management Science*, 62(8):2437–2455.
- Darr, E. D., Argote, L., and Epple, D. (1995). The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Science*, 41(11):1750–1762.
- den Boer, A. V., Meylahn, J. M., and Schinkel, M. P. (2026). Artificial collusion: Examining supracompetitive pricing by Q-Learning algorithms. *Management Science*.
- Denter, P. (2026). Strategic pricing and consumer welfare under one-sided price regulation. arXiv:2604.17576.
- Fisher, M., Gallino, S., and Li, J. (2018). Competition-based dynamic pricing in online retailing: A methodology validated with field experiments. *Management Science*, 64(6):2496–2514.
- Frondel, M., Thiel, P., and Vance, C. (2026). Bayesian causal estimation of Germany’s recent fuel price policy: The High Noon rule and the 2026 fuel tax discount. USAEE Working Paper 26-674, United States Association for Energy Economics.
- Gallego, G. and van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020.
- Harrington, Jr., J. E. (2022). The effect of outsourcing pricing algorithms on market competition. *Management Science*, 68(9):6889–6906.
- Heemels, W. P. M. H., Johansson, K. H., and Tabuada, P. (2012). An introduction to event-triggered and self-triggered control. In *Proceedings of the 51st IEEE Conference on Decision and Control*, pages 3270–3285.
- Jensen, R. J. and Szulanski, G. (2007). Template use and the effectiveness of knowledge transfer. *Management Science*, 53(11):1716–1730.
- Jung, L., Schildknecht, J., Gregor, L., and Haucap, J. (2026). Predictable prices, higher margins? early evidence on Germany’s 12 o’clock fuel regulation. ZEW–DICE analysis, ZEW – Leibniz Centre for European Economic Research and Düsseldorf Institute for Competition Economics.
- Kabul, M. O. and Parlaktürk, A. K. (2019). The value of commitments when selling to strategic consumers: A supply chain perspective. *Management Science*, 65(10):4754–4770.
- Kahl, M. P. and Wein, T. (2025). How to reach the paradise? Inside the Edgeworth cycle and why a gasoline station is the first to raise its price. *The Journal of Industrial Economics*, 73(2):351–383.
- Klenow, P. J. and Kryvtsov, O. (2008). State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *Quarterly Journal of Economics*, 123(3):863–904.
- Kraftstoffpreisanpassungsgesetz (2026). Kraftstoffpreisanpassungsgesetz vom 27. März 2026. BGBl. 2026 I Nr. 82.

- Levy, D., Bergen, M., Dutta, S., and Venable, R. (1997). The magnitude of menu costs: Direct evidence from large U.S. supermarket chains. *Quarterly Journal of Economics*, 112(3):791–824.
- Luco, F. (2019). Who benefits from information disclosure? The case of retail gasoline. *American Economic Journal: Microeconomics*, 11(2):277–305.
- Ma, W., Simchi-Levi, D., and Zhao, J. (2021). Dynamic pricing (and assortment) under a static calendar. *Management Science*, 67(4):2292–2313.
- Miklós-Thal, J. and Tucker, C. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4):1552–1561.
- Obradovits, M. (2014). Austrian-style gasoline price regulation: How it may backfire. *International Journal of Industrial Organization*, 32:33–45.
- Paarsch, H. J. and Schmedders, K. (2026). The reset premium: Option value under asymmetric price commitment. *Economics Letters*, 268:113146.
- Siemroth, C. (2026). The optimal reset-hour of a once-daily petrol price increase limit. arXiv:2603.18920.
- Tankerkoenig (2026). Historische Preisdaten. Official historical-data documentation. Official site capture retrieved May 8, 2026; the historical CSV repository is available to registered, verified users under CC BY-NC-SA 4.0 for noncommercial use, while the current API is separately licensed under CC BY 4.0.
- Wang, Z. (2009). (Mixed) strategy in oligopoly pricing: Evidence from gasoline price cycles before and under a timing regulation. *Journal of Political Economy*, 117(6):987–1030.
- Zbaracki, M. J., Ritson, M., Levy, D., Dutta, S., and Bergen, M. (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *Review of Economics and Statistics*, 86(2):514–533.
- Zipkin, P. H. (2000). *Foundations of Inventory Management*. Irwin/McGraw-Hill, Boston, MA.