

Pricing & revenue architecture

What is dynamic pricing?

Dynamic pricing sets the price of the same offer again as time, demand, capacity, cost or the buyer changes. In surveys, fairness turned on the reason.

Author	Published	Reading time	Words
Sinan Isoglu, MBA (Quantic) Commercial growth leader, lecturer and doctoral researcher	18 October 2026	17 min	3,670

Management summary Dynamic pricing is a policy under which the price of the same offer is set again as conditions change: the time, demand, remaining capacity or cost, the buyer's segment or the individual buyer. This concept guide defines it; separates its forms; distinguishes it from revenue management, price discrimination and price promotions; reports how, in a model estimated on one metals retailer's transactions, a price above a buyer's reference price went with a higher chance of moving to a more price-sensitive state; sets out which dynamic prices respondents judged unfair in Kahneman, Knetsch and Thaler's telephone surveys; and gives a one-page record for classifying a pricing rule. The sources are an operations model of revenue management (Maglaras and Meissner), a structural model of one seller's quotes and orders (Zhang, Netzer and Ansari) and those surveys.

Contents	§ 1	What is dynamic pricing?	3
	§ 2	What can a dynamic price respond to?	4
	§ 3	Is dynamic pricing the same as revenue management?	5
	§ 4	When is a dynamic price also price discrimination, or a promotion?	5
	§ 5	What did a study of one metals retailer find about prices and trust?	6
	Figure 1	Estimated moves between buyer states after a price below, at or above the reference	7
	§ 6	Which dynamic prices do people call unfair?	9
	§ 7	How can a team tell which kind of dynamic pricing it runs?	10
	Figure 2	A record for classifying a pricing rule	11
	§ 8	What do these sources leave open for software and subscriptions?	12

References

13

What is dynamic pricing?

§ 1 of 8

It is a price that is set again, for the same offer, whenever the conditions behind it change. The sources this guide uses come at it from two sides. In operations research, Maglaras and Meissner (2006)² study a firm that owns a fixed capacity of a resource and must use it within a finite horizon. In the first of their two formulations the firm is assumed “to have power to influence the demand for each product by varying its price”, and “the firm’s problem is to choose a dynamic pricing strategy for each of its products to optimize expected revenues”. In the second, prices are fixed and the firm chooses “a dynamic capacity allocation rule that controls when to accept new requests for each of these products”.

² Maglaras, C., & Meissner, J. (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management*, 8(2), 136–148.

In business markets the price can also vary across buyers. Zhang et al. (2014)³ observe that “In many B2B situations, sellers can easily vary prices across buyers and can even change prices between subsequent purchases of the same buyer.” They present their framework as an aid to sellers in “implementing first-degree and intertemporal price discrimination for long-run profitability”, and they contrast this with consumer retailers, which “are often limited in their ability to fully target prices for individual consumers because of logistical and ethical concerns”. They also name what makes it practical in business markets: “the reliance on salespeople, the need for product/service customization, and the volatility of commodity prices make targeting and intertemporal customization of prices feasible and desirable”.

³ Zhang, J. Z., Netzer, O., & Ansari, A. (2014). Dynamic targeted pricing in B2B relationships. *Marketing Science*, 33(3), 317–337.

In this guide, dynamic pricing means a policy under which the price of the same offer is set again when a condition the seller observes changes: the time, demand, remaining capacity or cost, the segment a buyer belongs to, or the individual buyer. The price is expected to change again when the condition does, which separates a dynamic price from a one-off revision of a list price. The policy need not be automated: the seller in Zhang et al.³’s study kept no price list and set each price or quote case by case. Dynamic pricing is distinct from revenue management, which can work with fixed prices, from a price promotion, which is temporary by announcement, and from price

discrimination, which it becomes when buyers pay different prices for the same offer.

What can a dynamic price respond to?

§ 2 of 8

Four kinds of condition, and one policy can combine them.

- **Time (time-based pricing).** The price changes with the date or with the time left before a deadline. In Maglaras and Meissner's² model the resources must be used within a finite horizon, and “capacity cannot be replenished up to that time”; one of their policies “reevaluates the fluid policy as a function of the current state and time-to-go”, that is, it solves a simplified version of the problem again from the capacity and time that remain.
- **Demand, capacity or cost (demand-based pricing).** The price changes with the state of the market or of the seller. In the same model the state is the remaining capacity; at the retailer Zhang et al.³ studied, the wholesale cost of the metal followed the London Metal Exchange daily. Surge pricing, a sharp rise when demand spikes, is a familiar version of this form.
- **Segment (segment-targeted pricing).** One price per group of buyers, which changes as buyers move between groups. Zhang et al. simulate a policy in which “only two optimal prices are determined”, “one price for each of the two states” in which their model places buyers.
- **The individual buyer.** A price for each buyer, possibly at each purchase. Their own proposal is to “compute optimal prices that are targeted for each transaction of each buyer so as to maximize long-term profits”. This is what is usually meant by personalized pricing.

Time and buyer are separate choices, and the label dynamic pricing is used for both. Zhang et al.³ keep them apart in the policies they compare: one in which “a single price is determined for each buyer for the entire 15-month planning horizon”, which is personalized but not dynamic over time; one that “chooses a single price for all buyers for the entire planning horizon”, which is neither, static pricing in the full sense; the segment policy; their proposal, which varies the price by buyer and over time; and a myopic pricing policy derived from it, in which “the seller maximizes profits only for the current purchase event, as opposed to the entire planning horizon”. In their counterfactual simulation the proposal earned the highest profit per buyer over nine months (Table 9, p. 333). The [essay on triggers](#),

constraints and explanations explains why a result of that kind is a model counterfactual, not a gain a seller can expect.

Is dynamic pricing the same as revenue management?

§ 3 of 8

No. Revenue management allocates a fixed capacity over a finite horizon, and changing prices is one way to do it. Maglaras and Meissner (2006)² write that such problems “gained interest in the late 1970s in the context of the airline industry” and have since been introduced in areas “such as hotels, cruise lines, rental cars, retail, etc.” Their second formulation, capacity control, solves the same problem with prices that “are assumed to be fixed either by the competition or through a higher-order optimization problem”, by deciding which requests to accept. A firm can therefore manage revenue without changing a price, and change prices with no scarce capacity to allocate. The comparison of revenue management and dynamic pricing sets out how to tell the two problems apart.

When is a dynamic price also price discrimination, or a promotion?

§ 4 of 8

Price discrimination means charging different buyers different prices for the same offer. A dynamic price becomes price discrimination when the condition it responds to is the buyer: a segment price or a personalized price. Zhang et al. (2014)³ use the term for their own policy, and their word intertemporal refers to the change over time, since the same buyer can pay a different price at the next purchase. A price that changes for everyone at once with cost or capacity treats all buyers at that moment alike, although it can still separate buyers by when they buy. Kahneman et al. (1986)¹ put the motive plainly: “A monopolist might attempt to increase profits by charging different customers as much as they are willing to pay.”

¹ Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, 76(4), 728–741.

A price promotion is a different object: a temporary price cut, announced, open to every buyer who meets its terms, with a stated or expected end. The end matters for how the next price is judged. In the surveys of Kahneman et al. (1986)¹, the most recent price serves as the reference “unless the terms of the previous transaction were explicitly temporary”. When a car model was in short supply, 71 per cent of respondents (N = 130) called it unfair for a dealer to sell at US\$200 above list price, while 42 per cent (N = 123) called it unfair for a dealer

that had been selling at US\$200 below list to return to list price, a significant difference. The authors read the gap as framing: the surcharge is a loss against an unambiguous list price, while the end of a discount can be coded as the elimination of a gain. The discount in that scenario had no stated end; the milder judgement rests on the list price, which at least some respondents appear to have kept as their reference. A promotion keeps its regular price in view in the same way. A dynamic price has no stated end and may have no regular price, so on this guide's reading each price it sets can become the buyer's next reference.

What did a study of one metals retailer find about prices and trust?

§ 5 of 8

Zhang et al. (2014)³ used 21 months of quotes and orders, January 2007 to September 2008, estimating their model on the first 18 and checking its predictions on the last three. The seller is what the abstract calls “a retailer in the industrial consumables domain” and the data section describes plainly: “Our data come from an East Coast aluminum retailer”. It cut metal to order for small and medium-sized industrial clients such as machine shops and fabricators. Because its cost followed the London Metal Exchange, “the seller does not maintain a price list” and set the price of each order or quote case by case. At each purchase event the buyer “either places a direct order, without asking for a price quote, or requests a price quote”, which it may then reject. The sample covers 1,859 buyers and 33,925 purchase events (Table 1, p. 320). The price and quantity quoted were “identical to the price and quantity reported on the final invoice for more than 99% of the orders”, so there was little negotiation.


From these transactions the model infers two latent states that buyers move between. “We therefore call State 1 the relaxed state and State 2 the vigilant state.” The model's probability of a quote request was 24 per cent in the relaxed state and 83 per cent in the vigilant state; the average absolute price elasticity across their four buying decisions was 1.4 against 3.1 (Table 5, p. 329). The authors call the states states of trust but add, of the difference between trust and a wider relationship quality, “Because the states are inferred from secondary data, we remain agnostic about this distinction.”

Each price entered the model as a gain or a loss against the buyer's reference price, the quantity-weighted average of the prices the buyer had seen at earlier purchase events. At the reference price, the authors note, “the likelihood of a buyer dropping from the relaxed to the vigilant state (14.3%) is almost twice as high as the likelihood of

moving from the vigilant to the relaxed state (7.7%)". Figure 1 shows the estimated transition probabilities after a price 10 per cent below, at or 10 per cent above the reference at the previous purchase event.

Figure 1 Estimated moves between buyer states after a price below, at or above the reference

In the estimated model, a relaxed buyer was vigilant at the next purchase event with a probability of 14.3% at the reference price and 21.1% after a price 10% above it.



Model estimates for one seller, not counts of observed moves: the states are inferred from quotes and orders, and the three price levels are values put into the model, not observed price changes. A buyer who does not move stays in its state, so each bar and the stay probability sum to 100 per cent.

Source Zhang et al. (2014), p. 330, Table 7: posterior mean of the transition matrix across buyers, from a hidden Markov model estimated on the first 18 months (January 2007 to June 2008) of transactions by 1,859 buyers of one East Coast aluminum retailer, with July to September 2008 held out for validation. Turns vigilant: a buyer in the relaxed state at one purchase event is in the vigilant state at the next. Turns relaxed: the reverse. The price is set against the buyer's reference price, the quantity-weighted average of past prices, at the previous purchase event. The table's probabilities are shown as percentages.

In the model, a price 10 per cent above the reference put the probability of turning vigilant at 21.1 per cent and that of turning relaxed at 4.3 per cent; a price 10 per cent below put them at 11.4 and 9.5 per cent. The authors conclude: "Thus, a price increase may have a long-term effect by transitioning the buyer to a (sticky) state of increased price sensitivity." Simulating one-off price changes, they also find that "the negative effects of a price hike (loss) on all four decisions persist longer than the positive effects of price drop (gain)".

The policy the authors derive from the model does not simply keep prices low to keep buyers relaxed. It “recommends a higher price for buyers in the relaxed state versus those that are in the vigilant state”, US\$3.63 against US\$3.20 per pound, because “Charging a higher price for buyers in the relaxed state can increase both the immediate profits and the reference prices in the long term.” The seller’s current policy, by contrast, “only mildly differentiates between buyers in the two latent states”. The simulation kept each buyer’s prices inside a range the buyer had already seen: “We choose to stay within each buyer’s historical range of experienced prices to avoid a drastic change in the price regime observed by each buyer”. In the simulation, the myopic version earned more in the first months and fell behind after about three, because, in the authors’ words, “charging lower prices results in lower reference prices and creates a downward pressure on the seller to continue offering lower prices”.

For swings in the metal price the model’s answer was asymmetric. After a 20 per cent rise in the exchange price, “the seller should increase the unit prices by 12% for buyers in the vigilant state but by only 4.6% for buyers in the relaxed state”; after a 20 per cent fall, the policy lowered prices by only 2.5 to 2.8 per cent. The authors call this consistent with the dual entitlement principle, citing Kahneman et al.¹ among others, while noting: “Although we do not model fairness directly, the external and internal asymmetric reference price effects in tandem with the latent trust state capture similar effects.”

The authors bound the result themselves. It is “a specific empirical application within the metal industry”; the data set “does not include competitive information”; the model assumes “that buyers are not forward looking with respect to the seller’s pricing decisions”; the results “are most suited to these more frequent buyers”, those with at least seven purchase events; and the model “needs to be extended to explicitly capture negotiation processes” where negotiation is common. The policy comparison is a simulation on 300 buyers over a nine-month holdout period, which the authors call “a heuristic for the optimal prices”.

Which dynamic prices do people call unfair?

§ 6 of 8

The classic evidence comes from telephone surveys. Kahneman et al. (1986)¹ read short scenarios, in which a merchant, landlord or employer changed a price or wage, to residents of Toronto and Vancouver between May 1984 and July 1985, and asked how fair the change was. They call the approach purely descriptive: a rule of fairness records what a substantial majority of respondents judged fair, not what is right. Their summary: “It is unfair to exploit shifts in demand by raising prices or cutting wages”, while it is acceptable “to raise prices (or cut wages) when profits are threatened and to maintain prices when costs diminish”. The principle behind it they call dual entitlement: “Transactors have an entitlement to the terms of the reference transaction and firms are entitled to their reference profit.”

Their scenarios touch each form described above.

- **Demand.** A hardware store that raised the price of snow shovels from US\$15 to US\$20 the morning after a snowstorm was judged unfair by 82 per cent of respondents (N = 107), and a grocer that raised the price of scarce apples by 25 per cent by 63 per cent (N = 102). In the authors’ words, “Community standards of fairness effectively require the firm to absorb an opportunity cost in the presence of excess demand, by charging less than the clearing price or paying more than the clearing wage.”
- **Time.** A peak-time surcharge drew the same objection: a majority judged it unfair for a popular restaurant to add a US\$5 surcharge for Saturday night reservations.
- **Segment.** A grocery chain that set prices 5 per cent higher in the one community where it had no competitor, with the same costs and sales volume, was judged unfair by 76 per cent (N = 101); versions with 10 and 15 per cent did not differ significantly.
- **The individual buyer.** A landlord who learned at renewal that the tenant had taken a job nearby, and so was unlikely to move, and who raised the rent US\$40 a month more than planned, was judged unfair by 91 per cent (N = 157). The authors conclude that “some forms of price discrimination are outrageous” and that “an action that deliberately exploits the special dependence of a particular individual is exceptionally offensive”.
- **Auction.** Selling the last doll in stock to the highest bidder a week before Christmas was judged unfair by 74 per cent (N = 101); saying that the proceeds would go to UNICEF brought that down to 21 per cent.

These are stated judgements about consumer, tenant and employee scenarios, collected in two Canadian cities in the mid-1980s: not observed purchases, and not business buying. Whether judgements of this kind change what buyers do is a separate question, which the price fairness guide takes up.

On this guide's reading, the two studies, read side by side, point to a tension that neither tests. The policy Zhang et al.³'s model favoured charged more to relaxed buyers, those who more often ordered without asking for a quote. In structure that resembles the landlord who charges more because the tenant is unlikely to move, a change 91 per cent of Kahneman et al.¹'s respondents judged unfair. The resemblance is partial. The landlord acted on an outside fact he had learned about one tenant. The relaxed state is inferred from the buyer's own ordering and can be left again; in the simulation its price stayed within the range each buyer had already seen, and the authors argue that the stickiness of that state and the higher reference prices act as a shield against buyers perceiving later prices as losses. The model's answer to commodity swings, which the authors describe as passing on "most of the cost increases" while keeping most of a saving, sits closer to what those respondents accepted. Whether business buyers who learned of such a policy would judge it the same way, the sources do not say.

How can a team tell which kind of dynamic pricing it runs?

§ 7 of 8

By answering six questions for each pricing rule: what the price changes with; whether buyers pay different prices at the same moment; whether the price is explicitly temporary and open to all; whether fixed capacity is being allocated before a deadline; which earlier price the buyer will compare it with; and what occasion the buyer will see for the change. The first four place the rule among dynamic pricing, price discrimination, a promotion and revenue management. The last two follow Kahneman et al.'s reference transaction and their finding that judgements turn on the occasion for a change.

Figure 2 A record for classifying a pricing rule

Write down what the price changes with, whether buyers pay different prices, whether it is temporary, whether capacity is allocated, the buyer's reference and the occasion before the rule runs.

Changes with	Across buyers	Stated end	Capacity	Reference	Occasion
Time, demand, capacity, cost, segment or buyer?	Different prices at the same moment?	Explicitly temporary and open to all?	Fixed capacity before a deadline?	Which earlier price will the buyer compare?	Cost, demand, market power or dependence?

Rising demand, market power and dependence were the occasions most respondents called unfair in the fairness surveys.

One row per rule. A yes in the second column makes the rule price discrimination, a yes in the third makes a price cut a promotion, and a yes in the fourth makes the rule part of a revenue management problem. The last two columns only flag a rule for review; the fairness record in the price fairness guide and the control map in the essay on triggers carry that review.

Source Author's decision worksheet. The first column follows the forms in Maglaras and Meissner (2006), pp. 136 to 139, and Zhang et al. (2014), pp. 317, 318 and 333; the capacity column follows Maglaras and Meissner (2006), p. 136; the reference and occasion columns follow Kahneman et al. (1986), pp. 728 to 735. The fields are reader inputs; no company data is supplied.

An illustration, with no company data: a B2B software firm renews annual contracts and proposes a renewal price for each account, higher where the account ran no competitive tender and did not involve its procurement team at the last renewal. The record classifies the rule as dynamic pricing by the individual buyer, and so as price discrimination; it has no stated end and allocates no capacity; the buyer will compare the new price with last year's; and the occasion, if the buyer learned it, would be its own low vigilance. That occasion is close to the dependence Kahneman et al.¹'s respondents condemned. It also resembles the relaxed state from which, in Zhang et al.³'s model, a price 10 per cent above the reference raised the probability of turning vigilant from 14.3 to 21.1 per cent, a risk their simulated policy accepted because the higher price still paid in the model. Written

down, the alternative is easier to see: tie the renewal price to a stated cost index, an occasion the buyer can check and one the surveyed public accepted for passing on costs. That is a claim about judgements, not about renewals. These are the author's examples and have not been tested.

What do these sources leave open for software and subscriptions? § 8 of 8

Each source has a narrow base. Maglaras and Meissner (2006)² analyse a model in which the firm “is either a monopolist or is assumed to operate in a market with imperfect competition” and cannot replenish capacity before the horizon ends, and they test their policies on numerical examples, not market data; in those examples the dynamic heuristics “tend to outperform the static one”, although one of them keeps prices fixed and uses capacity control, and the advantage of dynamic pricing was largest when capacity roughly matched demand at the revenue-maximizing prices and small when capacity was far above or below it. Zhang et al. (2014)³ estimate one seller's model and simulate the policies from it. Kahneman et al. (1986)¹ report judgements, not purchases.

On this guide's reading, one point carries over to software and subscriptions with some confidence, and two only as hypotheses to test. The first: name what the price changes with, because time, capacity, segment and buyer are different policies under one label. The hypotheses: that a price above a buyer's reference changes how the buyer buys next, not only whether, which rests on one seller's model; and that the occasion for a change decides how it is judged, which rests on consumer surveys. Two differences may matter. Most software has no fixed capacity to allocate, so the time and capacity forms that revenue management models study apply mainly to fixed resources such as implementation slots. And subscription buyers can negotiate and compare at renewal, while the buyers in Zhang et al.'s data rarely negotiated and the model assumes they did not anticipate prices. How either difference changes the results, the sources do not say.

 End of essay

References

- 1** Kahneman, D., Knetsch, J. L., & Thaler, R. (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *The American Economic Review*, 76(4), 728–741. <https://www.jstor.org/stable/1806070>
- 2** Maglaras, C., & Meissner, J. (2006). Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing & Service Operations Management*, 8(2), 136–148. <https://doi.org/10.1287/msom.1060.0105>
- 3** Zhang, J. Z., Netzer, O., & Ansari, A. (2014). Dynamic targeted pricing in B2B relationships. *Marketing Science*, 33(3), 317–337. <https://doi.org/10.1287/mksc.2013.0842>

How to cite this essay

Isoglu, S. (2026, October 18). What is dynamic pricing? Sinan Isoglu.
<https://isoglu.com/insights/journal/what-is-dynamic-pricing/>
Online version: <https://isoglu.com/insights/journal/what-is-dynamic-pricing/>

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher
© 2026 Sinan Isoglu · isoglu.com · All rights reserved