Original Research Paper

Volume - 16 | Issue - 08 | August - 2026 | PRINT ISSN No. 2249 - 555X | DOI : 10.36106/ijar

Economy

HOW LOW IS TOO LOW?

Saksham Agrawal

ABSTRACT While the conventional wisdom suggests that higher discounts will draw the attention of the consumers, it is often met with suspicion when the price drops beyond what they had expected. Such an outcome is referred to as the Cobra Effect of Price Discounts: a scenario where discounts discourage purchases from retailers. While the phenomenon is fact, this research paper aims to investigate its underlying causes: at what level of discount do consumers begin to perceive a product as of low quality?Though these discounting strategies may help retailers increase their sales, they can work to erode the trust that their consumers have of their retailers; thus, by determining such a threshold for discounts, retailers can use these discounts to their benefit.Using concepts from economics and asymmetric information, such as Akerlof's "market for lemons", Spence's signalling theory and more renowned ideas, we aim to not only determine the tipping point between a discount and a red flag for consumers, but also to test that assumption through a carefully controlled experiment. While much of the existing literature on the topic suggests that excessive discounts on products can lead to the erosion of the trust that consumers place in retailers, little is known about the exact tipping point for such an outcome and what factors may influence that tipping point.Thus, through this proposed research project, we aim to investigate the following question:At what discount depth does a discount on a product change from a signal of quality to one of low quality for the consumers — and does the type of product and justification for the discount play a role in that determination?

KEYWORDS :

INTRODUCTION

A paradox on the price tag

Think of two jars of olive oil, which are put next to each other; the first jar is priced at ₹500, whereas the second one is offered for only ₹120, which means 76 percent off from the initial price. According to the basic principles of demand theory, it is evident that the second one will sell quicker. The paradox is that people behave the opposite way.

They start pondering, drawing conclusions, saying "there is definitely something wrong about this", because the very feature that should attract attention turns out to repulse people instead.

This is not an experiment conducted on the market for olive oil; it is a phenomenon which emerges when there is no information provided regarding the quality of the goods. It is necessary to highlight that apart from pricing, there is additional information about the product and its market provided to customers on the label.

The Cobra effect

The Cobra Effect is a phenomenon used in economics to refer to a self-defeating process. Its origin dates back to the colonial era of Delhi, when the government, frightened by the menace of cobras, offered rewards to citizens for delivering every dead cobra, to solve the problem of an excess of these animals in the area. In order to earn money, citizens started breeding these snakes. When the program ended, the citizens released the now-worthless cobras out of anger, and there were even more cobras in the city after the end of the program (Siebert, 2001).

The reduction of prices produces a similar effect — a policy which aims at doing some good but harms the reputation of the products among consumers. This phenomenon is what we call the Cobra Effect of price discounts.

(We use the term as a memorable label, not in the strict sense — the classic Cobra effect breeds more of the bad outcome, whereas a discount erodes trust as an unintended side-effect.)

The gap this paper fills

That deep discounts can dent perceived quality is well established (we review the evidence in Section 2). What remains under-explored is this: at what level of discount does a bargain become a signal of low quality, and what determines that tipping point?

BACKGROUND AND RELATED WORK

Price as a signal of quality

The idea that buyers read quality from price is not a new one. Scitovszky (1944) first formalised that people judge a product based on its price when they lack information, arguing that it can even be the rational choice to some extent. Decades of experimenting and marketing research have backed this up.

A good example is a "social experiment" carried out by the Payless shoe store, called "Palessi", in 2018. They set up a fictitious luxury shop in Los Angeles, in which they placed ordinary shoes priced between $20 and $40. They then invited fashion celebrities to the opening ceremony, and these celebrities paid up to $645 for shoes that were nothing but ordinary Payless products.

Asymmetric information and the market for lemons

The basic principle behind this concept is asymmetric information. As per the "lemons theory" developed by Akerlof (1970), the seller knows the quality of the goods but the buyer does not; because he cannot tell the difference between a peach and a lemon, he will pay a price based on the average quality of the product. In case the average price is not high enough to support good products, the good products will exit the market.

Signalling theory

The signalling theory of Spence (1973) completes the equation. When a high-quality good is able to undertake an activity that cannot be profitably imitated by a low-quality good, that activity becomes a signal. An example is a warranty — only firms that are sure about their goods are able to issue long warranties; otherwise, the low-quality firm could not afford one.

Discount depth and perceived quality

Recent research closest to our interest, on discount depth, was conducted by Zheng and colleagues in the Journal of Retailing (2021). An inverted-U relationship between discount depth and perceived quality uncertainty was found. Moreover, experiments that vary discount levels — 20%, 40%, 60% and 80% — report the same: quality perceptions hold steady at small discounts and drop at bigger ones (IJFMR, 2025 — summary of the discount-framing literature). One such study in India, done in 2025 and titled "The Hidden Cost of Discounts", studies the impact of discount framing on consumer intention through surveys of consumers across different age groups and economic classes.

These studies help us build the base of our research, but our research differs from them in three respects: (i) a simple game-theoretic model with a prediction of a tipping point or threshold, (ii) an experiment to test the impact of having a reason for the discount on the threshold, and (iii) a comparison of inspectable and non-inspectable goods, a distinction borrowed from Nelson (1970).

Theoretical Framework: The Mathematics of Suspicion

The model consists of two halves that meet in the middle. The first part concerns the buyer's side — how a person who lacks information about the quality assumes it from the price. The next part conveys the seller's side. The two halves converge on a tipping point, and the experiment in Section 4 directly measures it.

The buyer's inference problem

There are two types of products — high quality (V_H) and low quality (V_L). Only the seller is aware of the true type. The buyer sees only the price and forms a belief (p = the probability the item is high quality). The highest value a rational buyer will pay is the expected value:
$EV = p \times V_H + (1-p) \times V_L$

For example, if the product is genuine V_H = ₹1,200 and if not V_L = ₹0. If the buyer, looking at the marketing and the seller, thinks it is genuine (p = 0.9), then EV = ₹1,080. If it is a coin-flip in the buyer's mind (p = 0.5), then EV = ₹600. Uncertainty alone nearly halves the worth of the product from the buyer's perspective.

The signalling game and the quality-assuring price

So far we have only seen the buyer's side; the full picture is a two-player signalling game: a seller who knows the quality choosing a price, and a buyer who sees the price and decides what to do. The seller would love it if every buyer believed the product was excellent. The obvious worry is a low-quality seller who simply charges a high price. What is stopping them?

- The trust comes from reputation and repeat business, and we can derive exactly how it builds up — setting up the seller's problem with some numbers:
- The cost of the high-quality product is cH = ₹450, and of the low-quality product is cL = ₹100.
- The buyer learns the true quality once he begins using the product, and through feedback and recommendations it becomes known to other potential customers. The penalty for selling low quality at a high price is the total loss of all future sales.
- The seller discounts future profits at δ = 0.9 (a future period's profit is valued at 90% of the current period's).

The honest seller receives the margin $(p - c_H)$ in each period, for a present value of $(p - c_H)/(1 - \delta)$. The dishonest seller earns the margin $(p - c_L)$ just once but is then discovered and earns no future profit. So it is worth staying honest — and charging a high price — only if:
$(p - c_H)/(1 - \delta) \geq (p - c_L)$

(here p = the price the seller charges). From this inequality we can calculate the quality-assuring price:
$p^* = [c_H - (1-\delta)c_L]/\delta = [450 - (0.1)(100)]/0.9 = ₹489$

This conclusion must be treated carefully, since it forms the basis of the entire analysis. If the price is below ₹489, it does not matter how honest the seller is — there is an incentive to cheat, and this explains why people find it hard to tell whether a cheap product is genuinely good. This value was calculated from the seller's incentives alone, without any added assumption. (For example, at p = ₹489, the value of honesty is (489 − 450)/0.1 = ₹390 while cheating gives 489 − 100 = ₹389, so honesty just wins; at ₹1,000 it is ₹5,500 > ₹900, so quality is credible; at ₹400 it is −₹500 < ₹300, so cheating takes over.)

The Bayesian backfire

Let the high-quality seller now make a steep price cut. The buyer faces a Bayesian question: given that I am seeing a really steep discount, what is the probability that the product is high quality? Bayes' rule answers this once supplied with likelihoods. For illustration — these are not measured probabilities — suppose steep price cuts are rare for high-quality sellers and common for lemon sellers liquidating stock:

- P(deep discount | High) = 0.2
- P(deep discount | Low) = 0.8

Then, starting from P(High) = 0.5, the buyer computes:
$P(\text{High} \mid \text{deep}) = (0.2 \times 0.5)/(0.2 \times 0.5 + 0.8 \times 0.5) = 0.10/0.50 = \mathbf{0.20}$
Given these likelihoods, the steep price cut lowers the buyer's probability from 50% to 20%. It is not the particular number that matters — that depends on the likelihoods used — but the principle: a price cut can reduce the perceived probability of quality, not only the price. This is how the Cobra effect operates. Whether actual buyers adjust their probabilities this way, and where, remains an open question — the one Section 4 addresses.

Locating the tipping point

Bringing it all together: as the discount rises, the asking price falls (good for the buyer), but the belief p falls as well (bad for the buyer). The buyer purchases only when EV is at least the asking price:
Buy ₹ $EV = p(d) \times V_H \geq$ asking price $= L(1-d)$
where d is the discount rate and L = ₹1,000 is the list price. At the tipping point, p falls so far that EV drops below the reduced asking price. Table 1 shows this with illustrative belief values chosen to reach a tipping point: the higher the discount, the lower the belief (slowly, then rapidly).

Table 1. A worked illustration of the mechanism, using illustrative belief values that the experiment estimates empirically.

Discount d	Asking price ₹1000(1−d)	Belief p = P(high)	EV = p × ₹1200	Buy? (EV vs price)
0%	₹1,000	0.90	₹1,080	Buy (+80)
20%	₹800	0.85	₹1,020	Buy (+220)
40%	₹600	0.75	₹900	Buy (+300) ₹ best deal
60%	₹400	0.30	₹360	DON'T BUY (−40) ₹ tips
80%	₹200	0.15	₹180	DON'T BUY (−20)

Reading top to bottom: from 0% to 40% the drop in perceived quality is slight, the price falls, and consumer surplus grows — a 40% discount is the best deal. Then, from 40% to 60%, perceived quality falls sharply (0.75 to 0.30): the offer enters the "too good to be true" zone. Expected value now falls faster than the price, so beyond 60% the offer is not worth taking — at 60% off the expected value is ₹360 while the price is ₹400. This is the interesting part, because in Section 3.2 we already pinned down the floor exactly.

Recall the quality-assuring price p* = ₹489. A discount cuts the price from L to L(1 − d). Quality stays credible only while the price stays above the floor, $L(1 - d) \geq p^*$. Rearranging gives the tipping-point discount:
$d^* = 1 - p^*/L = 1 - 489/1000 =$ **0.51 (about a 51% discount)**

So the model predicts a tipping point at about 51% off — squarely inside the 40–60% band the buyer-side table anticipated. It is not mere buyer panic; it is the exact point at which the seller's motivation to preserve quality disappears. Both lines of reasoning reach the same conclusion.

This also answers the second half of the question — what moves the tipping point. Because p* depends on the seller's patience and the cost of cheating, the threshold shifts in predictable ways (Table 2).

Table 2. How the tipping point moves with the seller's patience δ (list price ₹1,000, cH = ₹450, cL = ₹100).

Seller patience δ	Quality-assuring price p*	Tipping discount d*
0.80 (impatient / fly-by-night)	₹538	46%
0.90 (typical)	₹489	51%
0.95 (patient / repeat-business)	₹468	53%

The prediction is clear and testable. A seller who values the future — an established brand living on repeat business — can sustain quality at a lower price, and so can discount more deeply before buyers grow suspicious (53% before the red flag). A fly-by-night seller who barely values tomorrow loses credibility at a shallower discount (46%). Prediction: trusted brands should have deeper tipping points than unknown sellers — a 60%-off sticker that reassures on a household name may alarm on a name no one recognises.

The justification mechanism

Why does a festival "clearance" sale fail to raise suspicion? Because a justification changes the Bayesian arithmetic. If buyers assume that peaches and lemons alike discount during Diwali, then a deep discount is no longer diagnostic of quality: P(deep | High) and P(deep | Low) move closer together. Re-running Bayes with P(deep | High) = P(deep | Low) = 0.5 gives P(High | deep) = 0.5 — the belief does not move at all. The justification neutralises the signal, pushing the threshold deeper or removing it. In game-theoretic terms this is cheap talk that nonetheless works, because it supplies an innocent explanation for the low price.

Hypotheses

- **H1 (Tipping point).** Perceived quality and purchase intention stay relatively flat as the discount rises from 0% to around 40–60%, after which they drop sharply.
- **H2 (Justification).** Providing a credible reason for the discount pushes the tipping point to a deeper discount, or removes the drop in purchase intention.
- **H3 (Product type).** The tipping point occurs at a shallower discount for non-checkable products than for checkable products.
- **H4 (Introspection gap).** Consumers say they do not rely on price

to judge quality, yet behave as if they do — a gap between what they say and what they do.
- **H5 (Brand trust).** Following the sensitivity result in Section 3.4, the tipping point is deeper for a trusted brand than for an unknown one.

Experimental Design

Overview

The hypotheses will be tested with scenario-based experiments — the most popular, inexpensive and ethical approach for testing hypotheses about discount perception, and entirely feasible for a student using a paper-based or free online form.

Design

The experiment uses a 5 × 2 × 2 factorial design: discount level (0%, 20%, 40%, 60%, 80%) × justification (present / absent "festival clearance") × product type (checkable-a plain cotton T-shirt the consumer can see and feel; non-checkable-a phone charger or vitamins, whose quality is unknown before use). Discount level varies within subjects (each participant sees all discount levels for one of the two product types).

Participants

A convenience sample of students and friends, aiming for at least 20–30 participants per between-groups cell (a total of 80–120). The more, the better for a clean curve; but even a small sample should reveal an S-curve, because the discount effect is powerful.

Stimuli and procedure

1. The participant receives a credible list price for the product (e.g., a T-shirt "normally selling at ₹1,000").
2. The participant sees the identical product at all discount levels in a randomised sequence. In the justification condition, a banner reads "Festival Clearance — everything must go!"
3. At each level, the participant completes the measures below before moving on.
4. A debrief is conducted at the end (including the introspection-gap measure; see Section 4.5, Measure 4).

Measures

1. **Perceived product quality (1–7):** "How good do you think this product is?"
2. **Purchase intention (1–7):** "How likely are you to buy this product at this price?"
3. **Willingness to pay (₹) and a suspicion item (1–7):** "Do you suspect anything unusual about this deal?"
4. **Price-as-a-cue self-report (after debrief):** "Do you use price as a cue when judging a product's quality?" (Yes / No / Somewhat).

Analysis plan

1. **Locate the tipping point properly.** Plot mean perceived quality and purchase intention against discount depth and use segmented (piecewise) regression to find the breakpoint (e.g., a Davies test for a change in slope). The breakpoint, with a confidence interval, is the tipping point. A simple first-difference scan (the largest drop between adjacent levels) is used only as a descriptive check.
2. Compare the curves with and without justification (H2) and for checkable versus non-checkable products (H3). A rightward shift or flattening supports the hypothesis.
3. For H4, cross-tabulate stated cue reliance against revealed reliance (did their quality ratings actually move with price?).

Predicted Results and Analysis

The theory in Section 3 makes a specific, falsifiable prediction about the experimental findings, summarised in Table 3 and tested with the segmented regression of Section 4.6. The pattern to look for is a plateau followed by a sharp fall: perceived quality holding steady at shallow discounts and dropping once the discount passes the level at which a genuine seller would price. The figures below depict this expected pattern; the analysis then estimates the true breakpoint for our sample.

Table 3. Expected pattern of mean perceived quality (1–7) by discount depth, by product type and justification.

Discount	Checkable (no reason)	Non-checkable (no reason)	Non-checkable (+ reason)	Reading
0%	5.6	5.7	5.7	baseline trust
20%	5.6	5.5	5.6	deal, no alarm
40%	5.4	5.1	5.5	still fine
60%	4.9	3.2	5.0	non-checkable tips
80%	4.3	2.4	4.4	deep suspicion

Three predictions follow, one per hypothesis. First, perceived quality should stay roughly constant up to 40% and then fall sharply — an S-shaped curve, not a straight line — if a threshold exists at all. Second, the fall should be steeper for the non-checkable product (an experience good) than for the checkable one (a search good), because price is less informative once the buyer can see the T-shirt. Third, a justification ("festival clearance") should soften the fall for the non-checkable product, because it cancels the signalling weight of the price, as argued in Section 3.5.

Analysis. The breakpoint is estimated with the segmented regression of Section 4.6, with confidence intervals for each product and justification condition. The method is fixed before the data are analysed, so the breakpoint emerges from the data rather than from interpretation. The first-difference scan is reported only as an informal check.

DISCUSSION

Interpreting the tipping point

A confirmed threshold turns a familiar marketing rule of thumb into something measurable, with a mechanism behind it: past a certain point, the price is read not as a gift but as a confession. It matters not only whether a tipping point exists but what shape it takes. A plateau followed by a fall points to a signalling-style collapse, in which the buyer re-reads the price once it crosses the level a genuine seller would offer; a straight downward slope would instead be simple discounting with no inference at work.

Managerial implications

- **Stay under the line.** Discounts of 20–40% work well for non-checkable products without raising suspicion; discounts beyond 60% may cost more trust than they earn in traffic.
- **Justify deeper discounts.** A legitimate reason (clearance, festival, overstock, a new-brand launch) gives an innocent explanation that helps preserve faith in quality.
- **Add other signals for non-checkable goods.** When the buyer cannot judge quality before purchase, price does the heaviest signalling work, so discounts bite harder. Pair deep cuts with verifiable cues — certifications, reviews, warranties.
- **Price is not only money; it is also lost revenue**. When JCPenney adopted "everyday low prices" in 2012 and dropped its coupons and discounts, annual revenue fell about 25% — roughly $4.3 billion (widely reported; Forbes, 2017).
- **Why luxury never discounts.** Brands whose value is almost entirely a quality signal — designer goods, premium electronics — avoid discounts, because their tipping point is close to zero.

The say–do gap (an introspection gap)

If H4 holds, most participants will deny that price drives their judgement, even as their ratings track it closely. It is tempting to call this "preference falsification", but that would be inaccurate, and the difference matters. In Kuran's (1995) sense, preference falsification is when people conceal genuine preferences under social pressure. Our participants conceal nothing — they simply cannot see where their own judgement comes from. That is an introspection gap: the difficulty people have in reporting the true causes of their behaviour (Nisbett & Wilson, 1977). The methodological lesson is sharp: to study the psychology of price, watch behaviour, not statements about behaviour.

A note on what does not fit

It is worth marking the boundary. The effect studied here is the mirror image of Veblen goods, where desirability rises with price. That runs in the opposite direction and is not part of the mechanism examined in this paper.

LIMITATIONS AND FUTURE WORK

- **Hypothetical choices.** Scenario experiments measure intentions, not real spending; a field test with actual products and prices would strengthen external validity.
- **Sample.** A student sample limits generalisability; replication across ages and income levels (as in the 2025 Indian study) would help.
- **Single list price.** The threshold may itself depend on the price level and product category; future work can map thresholds across price tiers.
- **Repetition and reputation.** One-shot scenarios omit how repeated purchases and reputation reshape the signal.

CONTRIBUTIONS

Conceptual. The paper frames the price-cut backfire as a "Cobra effect

of price discounts", linking a popular marketing idea to the established theories of signalling and the market for lemons, and giving it a solid information-economics explanation.

Analytical. It presents a simple two-sided model — a buyer-side Bayesian inference and a seller-side reputation condition — that derives a quality-assuring price (the minimum price above which staying honest beats cheating) and, from it, a tipping point in discount depth (about 51% in our figures). It also shows how the threshold shifts with the seller's patience, yielding a testable prediction about trusted versus unknown brands. Simple mathematics produces numerical, falsifiable predictions.

Empirical. The study combines, in one experiment, the level of discount, its stated reason, and the search-versus-experience classification of goods (Nelson, 1970), measuring the threshold with a segmented-regression breakpoint. Analysing these three variables jointly, in an Indian consumer sample, is what makes the study original.

Taken together, the contribution is integrative: the paper combines three elements that have largely been treated separately — a mechanism, a model, and a measurement — and turns a loose notion ("don't discount too far") into a measurable question.

CONCLUSION

A price cut is meant to be good news. This paper has argued that, beyond a certain depth, it becomes bad news — because a price is not only what the buyer pays but also a message about what they are buying, and a number that looks too good to be true is read as exactly that. Bringing together the market for lemons, signalling theory, and a simple two-sided model, we located a tipping point at which a bargain turns into a red flag: in our worked figures, around 51% off, the point at which the price falls below the level that keeps a seller honest. We further argued that this point is not fixed — a credible justification pushes it deeper, non-checkable products reach it sooner, and trusted brands can discount further before suspicion sets in. These are stated as falsifiable predictions, to be confirmed or refuted by the experiment in Section 4. If they hold, the loose advice "don't discount too deeply" becomes a measurable line that retailers can locate and respect. How low is too low? Low enough that the buyer stops believing — and the aim of this paper is to find that point and explain why it is there.

Appendix A: Survey Instrument (template)

Repeat the block below for each discount level (0/20/40/60/80%), shown in randomised order. The justification condition adds the banner row at the top.

[Justification group only:] "Festival Clearance Sale — everything must go."
Product: _____ Normal price:₹_____ Today's price: ₹_____ (____% off)

1. How good do you think this product is?
   Very poor 1 — 2 — 3 — 4 — 5 — 6 — 7 Excellent
2. How likely are you to buy it at this price?
   Very unlikely 1 — 2 — 3 — 4 — 5 — 6 — 7 Very likely
3. What is the most you would pay for it?
   ₹_____
4. Does anything about this offer feel "off" or too good to be true?
   Not at all 1 — 2 — 3 — 4 — 5 — 6 — 7 Very much

Debrief (asked once, at the end):

1. In general, do you judge a product's quality by its price?
   Yes / No / Somewhat

For analysis, match each person's debrief answer against whether their quality ratings actually shifted across the discount ladder. A "No" paired with ratings that clearly moved with price is the introspection-gap signature (H4).

REFERENCES

Note: references with a hyperlink were verified via online search for this draft. Confirm the volume, issue, and page numbers against the source material before final submission.

1. Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3), 488–500.
2. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263–291.
3. Kuran, T. (1995). Private Truths, Public Lies: The Social Consequences of Preference Falsification. Harvard University Press.
4. Nelson, P. (1970). Information and consumer behavior. Journal of Political Economy, 78(2), 311–329.
5. Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. Psychological Review, 84(3), 231–259.
6. Rao, A. R., & Monroe, K. B. (1989). The effect of price, brand name, and store name on buyers' perceptions of product quality. Journal of Marketing Research, 26(3), 351–357.
7. Scitovszky, T. (1944). Some consequences of the habit of judging quality by price. Review of Economic Studies, 12(2), 100–105.
8. Siebert, H. (2001). Der Kobra-Effekt: Wie man Irrwege der Wirtschaftspolitik vermeidet. Deutsche Verlags-Anstalt. (Origin of the popular "Cobra effect" term.)
9. Spence, M. (1973). Job market signaling. Quarterly Journal of Economics, 87(3), 355–374.
10. Thaler, R. (1985). Mental accounting and consumer choice. Marketing Science, 4(3), 199–214.
11. Zeithaml, V. A. (1988). Consumer perceptions of price, quality, and value: A means-end model and synthesis of evidence. Journal of Marketing, 52(3), 2–22.
12. Zheng, et al. (2021). Retail price discount depth and perceived quality uncertainty. Journal of Retailing. ScienceDirect link
13. Heinsalu, S. (2018). Competitive pricing despite search costs if lower price signals quality. Working paper. arXiv link
14. "The Hidden Cost of Discounts: Do Consumers Value What They Pay Less For?" (2025). International Journal for Multidisciplinary Research, 7(1). IJFMR link
15. MIT thesis (n.d.). Pricing with quality perception: Theory and experiment. MIT DSpace. DSpace link
16. Mourdoukoutas, P. (2017). A strategic mistake that still haunts JC Penney. Forbes. (JCPenney 2012 pricing case; ~25% revenue decline.) Forbes link
17. CBS News (2018). Payless sold discount shoes at luxury prices — and it worked. (The "Palessi" experiment; markups up to 1,800%.) CBS News link