

THE CONFLICTED ALGORITHM. WHO AI POWERED BEAUTY RECOMMENDATIONS REALLY WORK FOR, STUDIED THROUGH NYKAA'S SKIN SCAN

Asavari Agrawal

Inventure Academy, Bangalore, India

DOI: 10.46609/IJSSER.2026.v11i08.024 URL: <https://doi.org/10.46609/IJSSER.2026.v11i08.024>

Received: 10 August 2026 / Accepted: 21 August 2026 / Published: 24 August 2026

ABSTRACT

Beauty shopping in India has moved online at remarkable speed, and the platforms that lead it now use artificial intelligence to tell people what to buy. In November 2025 Nykaa launched Skin Scan, a tool that turns a selfie into what the company calls a personalised skin consultation and then recommends products. This paper asks a question the marketing does not answer. When the company that runs the AI advisor also owns many of the brands it could recommend and can earn a higher margin on those owned brands, whose interest does the recommendation serve. Drawing on published financial data, I show that Nykaa has a measurable reason to favour its own products, since third-party analysis estimates materially higher margins on owned brands than on comparable third-party products, though Nykaa does not disclose audited product-level margins. I argue this is a conflict of interest of the kind competition regulators have examined in marketplaces such as Amazon, but with a sharper edge, because a skin diagnostic borrows the authority of an expert and is harder to question. I then set out an original method, a paired scoring audit built around a Consumer Optimality Score and a Platform Incentive Score, that lets a researcher test whether these recommendations follow consumer value, platform value, or both. The paper contributes a framework and research design rather than a verdict and argues that the conflict is structural to vertically integrated beauty platforms rather than the fault of any single firm.

Keywords: artificial intelligence, self-preferencing, information asymmetry, private label, gross margin, beauty e-commerce, Nykaa

1. Introduction

It was close to midnight, and I was lying in bed with my phone when the app offered to look at my skin. I had been scrolling through serums I did not need, the way you do when you are meant

to be asleep, and a bright card appeared inviting me to try Skin Scan. I took the selfie. A few seconds later the screen filled with a tidy verdict about my pores and my hydration, and beneath it sat three products I should buy, each matched neatly to a concern I did not know I had until that moment.

It felt helpful. It also felt like advice, the kind you get from a pharmacist or a dermatologist, someone whose job is to know more than you and to be on your side. That was the feeling that stopped me. I realised I had no way of telling whether the app was recommending what was best for my skin or what was best for its own business. The recommendation looked personal and neutral. Behind it was a company that sells products, some of which form part of its owned-brand portfolio and can contribute disproportionately to gross-margin expansion.

That gap, between how a recommendation feels and what drives it, is the subject of this paper. I want to be careful from the start. I do not claim that Nykaa deceives anyone, and I have not caught the company doing anything wrong. What I set out to do is show, using its own published numbers, that the economic incentive to favour its own products is potentially significant, to explain with standard tools from economics and accounting why a shopper cannot easily notice that steering, and then to propose a way to measure whether it happens. If a human shop assistant earned double the commission on one brand and pushed it without saying so, we would call that a conflict of interest. The question here is what we should call it when the assistant is an algorithm the shop built itself.

The paper proceeds as follows. Section 2 sets out the market and its competitive structure. Section 3 examines the economics of Nykaa's owned brands using its published accounts. Section 4 develops a theoretical framework for conflicts of interest in algorithmic recommendation. Section 5 describes Skin Scan and the question it raises. Section 6 presents the proposed audit methodology, Section 7 illustrates it, and Sections 8 to 10 discuss implications, limitations and conclusions.

2. The Indian Beauty Market and Its Competitive Structure

India's beauty and personal care sector is one of the country's largest consumer markets, and although online shopping still accounts for a minority of it, that share has been rising quickly. A large market with low but rising online penetration is exactly the setup that attracts heavy investment, because there is so much room to grow. That is what happened.

Nykaa began this shift. Falguni Nayar founded it in 2012, selling cosmetics from a single Mumbai office after a long career in banking. The company grew by giving shoppers in smaller cities access to brands they could not otherwise reach, and by wrapping that access in reviews, tutorials and content so that buying felt guided rather than random. By late 2024 it reported

roughly a thirty percent share of the beauty vertical, a cumulative base of around thirty-seven million customers, and more than two hundred stores [1]. Its financial performance has kept improving. In the financial year 2026, revenue from operations rose about twenty-six percent to 10,022 crore rupees, and full year EBITDA climbed fifty-nine percent to 752 crore rupees, lifting the EBITDA margin to 7.5 percent from six percent a year earlier [2]. Profit after tax reached 204 crore rupees, up 183 percent [2, 3].

Nykaa no longer has the field to itself. Tira, backed by Reliance Retail, entered with an offline first approach and opened eleven stores within about a year of launching, using the group's existing retail muscle [4]. Fashion platforms such as Myntra pushed into beauty as well, and beneath all of them quick commerce apps began delivering everyday products in minutes. Analysts have been direct about what this means for Nykaa. Karan Taurani of Elara Capital has argued that the realistic goal for Nykaa is to hold its share of around twenty-five to thirty percent rather than keep gaining it, because vertical players, super apps and quick commerce are now all competing for the same shopper [4].

In the language of economics, the premium online beauty market is increasingly concentrated among a small number of well-funded platforms. These firms compete on branding, exclusivity and increasingly on technology rather than on price alone. Barriers to entry are high, because trust, brand relationships and data all take years to build, which is why the contest is between a handful of well-funded platforms and not a crowd of small ones. That structure matters for everything that follows, because market power is what turns a small nudge inside a recommendation into a real shift in sales.

3. The Economics of Owned Brands

To understand the temptation to steer, it helps to look at how Nykaa earns its gross profit, and here the accounts speak plainly. Unlike a pure marketplace, Nykaa's beauty and personal-care business is predominantly inventory-led. It sources products from brands and authorised distributors and sells them on its own account. This inventory-led model gives Nykaa substantial control over the products it sells and provides the commercial context for its owned-brand strategy.

The centre of that strategy is House of Nykaa, Nykaa's portfolio of owned beauty and fashion brands, reported as twelve brands across beauty and fashion in the financial year 2026 [8]. These include Nykaa Cosmetics, Kay Beauty, Dot and Key and Earth Rhythm. When Nykaa sells an owned-brand product, it participates in both the brand economics and the retail economics, creating the potential for a higher contribution margin than on a comparable third-party product. Third-party industry analysis has estimated materially higher gross margins for Nykaa's owned

products than for comparable third-party products, including an illustrative comparison of approximately forty-four percent for an outside lipstick against sixty to seventy percent for a Nykaa-owned lipstick [5]. These are industry estimates rather than audited product-level margins disclosed by Nykaa, so they should be treated as evidence of the direction of the incentive rather than its exact magnitude. The economic implication is nevertheless clear. For the same rupee of sales, an owned brand can potentially leave Nykaa with more gross profit than a comparable third-party brand.

This is why the owned brands have scaled so hard. By the third quarter of the financial year 2026 the House of Nykaa beauty portfolio delivered 775 crore rupees of gross merchandise value in a single quarter, reaching an annualised run rate near 3,100 crore rupees, with Dot and Key at roughly a 1,900-crore rupee run rate, Kay Beauty near 500 crore rupees and Nykaa Cosmetics around 480 crore rupees [6]. It shows up at the top of the income statement too. The rising share of owned brands, together with operating leverage, lifted Nykaa's gross margin to 45.1 percent in the financial year 2026, and the company itself credits the improvement to the shift in its sales mix toward owned brands [7, 3].

Two points follow. First, this is not only a Nykaa habit. Business Standard reported in June 2026 that both Tira and Nykaa are scaling private label brands precisely because the margins are better and give more control over pricing [8]. The incentive runs across the whole industry. Second, and this is where the accounts meet the ethics, a business that can earn more on its own products has a persistent economic incentive to increase the likelihood that those products are chosen. In a shop with human staff, we would name that risk without hesitation. The question this paper raises is what changes when the one doing the recommending is an algorithm the same company designed.

4. A Theoretical Framework for Algorithmic Conflicts of Interest

Economics offers a precise way to describe the danger. A recommendation system acts as an agent for the shopper, who is the principal. The shopper wants the best product for their skin and their budget. The catch is that the agent is not paid by the shopper. It is built and funded by the platform, whose interests only partly line up with the shopper's. This is the classic principal and agent problem, and it opens the door to moral hazard, where the better-informed side acts in its own interest because the other side cannot easily observe what it is doing.

What makes observation so hard is information asymmetry. The platform knows the margin on every product, the promotion deals attached to them, and the exact logic of its own tool. The shopper knows almost none of this. Skincare adds a further twist. Many of these products are what economists call credence goods, meaning the buyer struggles to judge how well they

worked even after using them, because skin changes slowly and for many reasons at once. If you cannot fully verify quality even after you have bought and used something, you are unusually dependent on whoever advised you. That dependence is valuable to the advisor, and it is exactly where a conflict of interest can hide.

This is not a concern invented for the occasion. Competition authorities have taken the general version of it seriously. The practice of a platform favouring its own products in its own rankings is called self-preferencing, and it sits at the centre of real cases, including the European Commission's investigation into Amazon's Buy Box and the OECD's review of algorithmic competition [9]. Independent researchers have used algorithmic audits to investigate whether Amazon's systems favour its own products, documenting elevated visibility for its private-label items in search and recommendation placements [10]. Nykaa's own structure invites the same question. As one industry analysis put it, the company is on both sides of its own shelf, since it decides what shows up in search, how often a brand is shown, and what its advisors recommend [7].

The reason this deserves fresh study, rather than being treated as settled, is the particular form the advisor now takes. A plain row of suggested products is easy to read as advertising. A skin diagnostic is different. It arrives with the authority of expertise, closer to a consultation than to a shelf display, which is exactly how Nykaa presents Skin Scan [11]. When advice sounds like a professional opinion, a shopper lowers their guard. A conflict of interest is always strongest when the person receiving the advice does not realise there is one.

5. The Case of Nykaa's Skin Scan

In November 2025 Nykaa introduced Skin Scan on its app. The tool guides a user to take a selfie, analyses it, and recommends a set of products matched to the skin concerns it identifies. Nykaa presents it as bringing the clarity of an in-store consultation online, and describes it as designed to analyse patterns across customer concerns and deliver personalised recommendations that improve conversion and deepen engagement [11].

Read closely, that description holds both halves of the tension examined here. Skin Scan is genuinely useful, and it is openly designed to improve conversion, which is a commercial goal. Those two things are not the same, and they are not always opposed either. A recommendation can be good for the shopper and good for the platform at once, which is the happy case. It can also be good for the platform and only fine for the shopper, which is the outcome a conflict of interest tends to produce over time. The marketing cannot tell us, which is happening, because the marketing has a stake in the answer. The only way to know is to measure it and making that measurement possible is the purpose of the method that follows.

6. Proposed Methodology, the Consumer Optimality and Platform Incentive Scores

The main contribution of this paper is a way to test the question rather than argue about it. The idea is to score two different things for every product the AI could recommend, and then to see which of the two the actual recommendations follow.

The first is the Consumer Optimality Score, or COS. It measures how good a product is for the shopper, given their stated concern, and it combines four things a shopper could in principle observe. The first is concern suitability, meaning how well the product's stated purpose and target skin type match the situation. The second is ingredient and formulation fit, meaning whether the product's formulation and active ingredients are supported as appropriate for the stated concern, taking account of formulation context rather than simply counting ingredients. Examples may include salicylic acid for acne, niacinamide for oiliness, humectants such as hyaluronic acid for hydration, and retinoids for photoageing. The third is rating quality, measured as the star rating weighted by how many reviews stand behind it. The fourth is price to the consumer, where a lower price counts as better value, all else equal. With each component normalised to a zero to one range within the choice set, the score is

$$COS = w1 \text{ (suitability)} + w2 \text{ (ingredient fit)} + w3 \text{ (rating quality)} + w4 \text{ (1 minus normalised price)}.$$

The second is the Platform Incentive Score, or PIS. It measures how much the platform gains if a product is chosen. This is the harder score to build honestly, because the real drivers, the exact margins and advertising deals, are internal to Nykaa and cannot be seen from outside. Rather than pretend to know them, the method uses only observable stand ins and says so plainly. The first component is margin, stood in for by whether the product is an owned House of Nykaa brand, which is defensible because third-party analysis estimates a materially higher margin on owned brands, on the order of sixty to seventy percent against about forty-four percent on comparable third-party products, even though Nykaa does not disclose audited product-level margins [5]. The second is promotion, stood in for by whether the product carries a sponsored, promoted or bestseller label. The third is transaction value, stood in for by price, since a higher priced sale is worth more to the platform. The fourth is strategic value, stood in for by whether the product is exclusive to the platform or newly launched, since platforms are known to push both. The score is

$$PIS = v1 \text{ (owned brand)} + v2 \text{ (promoted)} + v3 \text{ (normalised price)} + v4 \text{ (strategic)}.$$

A safeguard is built into this that matters for credibility. The two scores share price, which enters them with opposite signs because price can create a genuine tension between consumer value and platform revenue. Other components are conceptually distinct, although they may still be

empirically correlated. As a robustness check, the audit should report correlations among the score components and repeat the analysis with price removed from one of the scores.

The procedure runs in five steps. First, profile how Skin Scan takes input, since whether it can be steered by declared concerns or only reads the selfie changes the design. Second, build thirty to fifty scenarios that cross common skin concerns with skin types. Third, for each scenario, record the product the AI recommends and assemble a comparison set of five to eight genuinely comparable alternatives using a predefined sampling rule. Where feasible, the eligible universe should be established independently of the Skin Scan results before the recommendation is observed, to reduce selection bias and avoid allowing the algorithm to define its own benchmark. Fourth, score every product in every set on COS and PIS, ideally with a second person independently scoring the judgement-based components so that inter rater agreement can be reported. Fifth, analyse.

The analysis answers three questions. Does the recommended product tend to be the highest COS option in its set, which is what a purely consumer serving tool would produce. Are owned brands recommended more often than their simple availability in the comparison sets would predict, which is the clearest single sign of steering and mirrors the Amazon audit logic. And, in a logistic regression of whether a product is recommended on its COS and PIS, does PIS still predict recommendation once COS is held constant. That last test is the strongest in the proposed design, because controlling for measured consumer value helps distinguish product suitability from platform-linked recommendation patterns. If more platform-favourable products are more likely to be recommended even after measured consumer value is controlled for, that would be evidence consistent with a conflict of interest expressing itself inside the algorithm, although it would not by itself establish causality.

To keep the weights from being accused of manufacturing a result, the baseline treats all weights as equal, and the study then repeats the analysis under several weighting schemes to show that the direction of the finding holds. Robustness to weighting is itself a result worth reporting.

7. An Illustrative Application

The example below uses invented products and invented numbers, only to show how the instrument works. It is not data, and it is labelled clearly so that no reader mistakes it for a finding.

Consider a scenario for oily skin with mild acne. The AI recommends Product A. The comparison set also holds Products B, C and D, all serums aimed at acne and oiliness. Product A is an owned brand, priced high, carrying a bestseller tag, containing niacinamide but no salicylic acid, with a good rating on a moderate number of reviews. Product B is an outside brand,

cheaper, containing both salicylic acid and niacinamide, with a strong rating on a very large number of reviews. On the consumer side, Product B looks stronger, because it has better ingredient fit for acne, a lower price and more credible ratings, so its COS would come out higher than Product A's. On the platform side, Product A looks stronger, because it is an owned brand, carries a promotion and sells at a higher price, so its PIS would come out higher.

In this made-up case the AI recommended the higher PIS product rather than the higher COS one. If that happened once it would mean nothing, since any single recommendation can be defended. The power of the method is that it repeats the comparison across thirty to fifty scenarios and asks whether the pattern is systematic. One instance is an anecdote. A steady tilt across dozens of independent scenarios, holding consumer quality constant, is evidence.

8. Discussion and Implications

Suppose an audit of this kind found that owned brands were recommended far more often than their share of comparable options would justify, and that PIS predicted recommendation even after controlling for COS. In welfare terms that would suggest a quiet transfer of value from consumers to the platform. Shoppers would be paying a little more, or getting a slightly less suitable product, in exchange for advice they believed was neutral. That is a small loss for each individual and potentially a very large one in aggregate, given that it operates on a platform serving tens of millions of beauty customers. Many modern consumer harms have exactly this shape, minor for the individual and significant at scale.

The obvious objection is that competition should fix this. If Nykaa steers people toward worse value, surely Tira or Purplle will win them over with better advice. There are two reasons to doubt that comfort. The first is that the incentive is not unique to Nykaa. Because owned brands carry higher margins for everyone, every vertically integrated platform feels the same pull, and Business Standard has confirmed that Tira is scaling private labels for the very same reason [8]. When rivals share an incentive, competition can settle on it rather than compete it away. The second reason is the credence goods problem. If shoppers cannot easily tell which platform gives better skincare advice, they cannot reward the better one, and the disciplining force of competition weakens. This is why the conflict is best framed as structural. It is a feature of how vertically integrated, AI led beauty platforms are built, not a lapse by one company that a rival will automatically punish.

None of this requires bad intent. A recommendation engine tuned to raise conversion and revenue, trained on a catalogue where owned brands are more profitable, can produce self-preferencing simply as a byproduct of doing its job well. That is what makes this an economics and governance question rather than a morality tale. If a remedy is needed, it looks less like

blaming a firm and more like the tools economists already reach for, meaning transparency about how recommendations are generated, clear labelling of owned brands, and independent auditing of the kind this method is designed to enable.

The implications differ by stakeholder. For consumers, the practical lesson is that an AI recommendation is not automatically neutral, and the information needed to judge it, namely margin and promotion status, is exactly the information a platform does not display. For Nykaa, there is a genuine trade-off between the short-term margin gain from owned brands and the long term value of being believed, since trust in neutral advice is part of why customers come at all. For a competitor such as Tira or Purpille, the question is whether honest recommendation can be a real advantage when a shopper cannot verify honesty, and how such a claim could ever be proven. For a regulator, self-preferencing has been examined in general marketplaces, but a diagnostic that borrows the authority of expertise may warrant stricter disclosure than an ordinary sponsored listing.

9. Limitations

The limits of this study should be stated openly. The Platform Incentive Score relies on public stand ins rather than Nykaa's real internal figures, so it captures the direction of the incentive rather than its exact size. The scoring of suitability and ingredient fit involves judgement, which is why independent second scoring and a reported agreement measure are part of the design. Any real audit would use a modest sample and would be a snapshot in time, since prices, promotions and the tool itself keep changing. Recommendations can also reflect genuine popularity, inventory availability, advertising, historical conversion, brand exclusivity or other factors rather than steering. Controlling for measured consumer quality is therefore necessary but may not be sufficient for causal identification. Framed as an exploratory audit, these are boundaries to state clearly, not weaknesses to conceal.

10. Conclusion

The Indian beauty market grew by giving people access and advice they felt they could trust. Artificial intelligence now delivers a growing share of that advice, and the companies that own the advice also own a growing share of the products it points to. That overlap creates a conflict of interest that is easy to feel and hard to see, made harder because a skin diagnostic speaks in the voice of an expert. This paper has used published accounts to show that the economic incentive to favour owned products is potentially significant, explained with standard economic ideas why a shopper cannot easily detect it, and offered a concrete method for measuring whether steering happens. It stops short of a verdict on Nykaa, on purpose, because the responsible next step is measurement and not accusation. The wider point is the one that matters

most. As recommendation moves from human staff to algorithms owned by the seller, the old idea of a conflict of interest does not disappear. It simply becomes invisible, and finding ways to see it again is one of the real tasks of studying markets in the age of AI.

References

- [1] Forbes India (2024) How Nykaa is taking on intensifying competition in the beauty industry, 11 December. Available at <https://www.forbesindia.com/article/billionaires/how-nykaa-is-taking-on-intensifying-competition-in-the-beauty-industry/94824/1> (accessed August 2026).
- [2] Forbes India (2026) W-Power 2026, Falguni Nayar, the woman who redefined beauty retail in India, 22 June. Available at <https://www.forbesindia.com/article/w-power-2026/falguni-nayars-nykaa-is-much-more-than-a-marketplace/2995203/1> (accessed August 2026).
- [3] Markhub24 (2026) Nykaa's Private Label Expansion Strategy, 8 July. Available at <https://www.markhub24.com/post/nykaa-s-private-label-expansion-strategy> (accessed August 2026).
- [4] The Core (2024) Can Nykaa Keep Its Market Crown in This Beauty Battle, 29 August. Available at <https://www.thecore.in/business/nykaa-tira-myntra-beauty-retail-competition-754285> (accessed August 2026).
- [5] Acheron Research (2026) Nykaa, India's Beauty Platform Playbook. Available at <https://acheronresearch.substack.com/p/nykaa> (accessed August 2026).
- [6] BeautyMatter (2026) Nykaa and the Rise of India's Beauty Economy. Available at <https://beautymatter.com/articles/how-nykaa-built-indias-beauty-market> (accessed August 2026).
- [7] Inc42 (2026) Inside Nykaa's Inventory Model, Where the Margin Comes From. Available at <https://inc42.com/features/inside-nykaa-inventory-model-where-the-margin-comes-from/> (accessed August 2026).
- [8] Business Standard (2026) Online beauty marketplaces push for growth with their in-house brands, 21 June. Available at https://www.business-standard.com/industry/news/beauty-online-market-places-push-for-growth-with-their-in-house-brands-126062100529_1.html (accessed August 2026).

- [9] OECD (2023) Algorithmic Competition, Competition Policy Roundtable Background Note. Available at https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/05/algorithmic-competition_2be02d00/cb3b2075-en.pdf (accessed August 2026).
- [10] Dash A, Chakraborty A, Ghosh S, Mukherjee A, Frankenreiter J, Bechtold S, Gummadi KP (2024) Antitrust, Amazon, and Algorithmic Auditing. arXiv preprint arXiv:2403.18623. Available at <https://arxiv.org/abs/2403.18623> (accessed August 2026).
- [11] Media Infoline (2025) Nykaa launches Skin Scan, an AI led skin diagnostic experience, 18 November. Available at <https://www.mediainfoline.com/advertising/nykaa-launches-skin-scan-an-ai-led-skin-diagnostic-experience> (accessed August 2026).

Appendix A. The Audit Instrument

The instrument below is the scoring sheet designed for the proposed Skin Scan audit. Each product that appears in a scenario, whether it was recommended or not, is scored on every field. Components are normalised to a zero to one range within each choice set before the scores are combined.

Table 1. Scoring fields for the proposed Skin Scan audit.

Field	What is recorded	Feeds into
Scenario ID	Concern and skin type combination	Both
Product and brand	As shown in the app	Both
Recommended	Whether the AI recommended it, yes or no	Outcome
Concern suitability	Match between product purpose and scenario, scored 1, 0.5 or 0	COS
Ingredient fit	Appropriate active ingredients present and how prominent	COS
Rating quality	Star rating weighted by number of reviews	COS
Price	Selling price, and price per ml or gram where possible	COS and PIS
Owned brand	Is it a House of Nykaa brand, yes, no, or part owned	PIS
Promoted	Sponsored, promoted or bestseller tag present	PIS
Strategic	Exclusive to platform or newly launched	PIS

Note. COS denotes the Consumer Optimality Score and PIS the Platform Incentive Score. The instrument is designed by the author and does not reproduce any external data source.