Algorithmic Illusion? An Empirical Study of Generative AI Image Ban on a Food Delivery Platform

Jingwei Dai[1], Beibei Li[1], and Rahul Telang[1]

[1]Heinz College of Information Systems and Public Policy, Carnegie Mellon University,
Emails: jingweid@andrew.cmu.edu, beibeili@cmu.edu, rtelang@cmu.edu

December 2, 2025

Abstract

Platforms increasingly use generative AI to create and enhance content, but it remains unclear whether such images boost user engagement or erode consumer trust. We investigate this question in the context of India's food delivery market. Following the surge of generative AI technologies, India's two largest food delivery platforms—Zomato and Swiggy—began offering generative AI tools to help restaurants beautify dish images starting in late 2023. However, in August 2024, Zomato enacted a platform-wide ban on AI-generated or AI-modified images, citing concerns over authenticity, while its main rival, Swiggy, continued to allow their use. Using a unique panel of consumers' purchases across both platforms from March 2024 to February 2025, we implement a difference-in-difference design with Zomato restaurants as the treated group and Swiggy restaurants as the control group. We find that, following the ban, Zomato restaurants experienced significantly greater increases in order volume, revenue, and customer counts, compared with restaurants on Swiggy. Consumers also became more willing to try previously unvisited restaurants on Zomato. Multi-modal image forensic analyses and three state-of-the-art generative AI detectors consistently show that Swiggy's dish images were more likely to be AI-generated after the policy divergence. Moreover, Swiggy restaurants with more "AI-like" images experienced larger sales declines compared with Zomato restaurants, aligning with anecdotal evidence that consumers dislike AI-modified dish images. Overall, our findings suggest that banning generative AI-modified dish images may buoy merchant sales and enhance platform competitiveness. These results highlight the need for platforms and merchants to exercise caution when deploying generative AI in contexts where content authenticity is critical to consumer decision-making.

Keywords Generative AI, food delivery, platform competition, algorithm aversion, image forensics

1 Introduction

The rapid development of generative artificial intelligence (Generative AI) has created both opportunities and challenges for businesses and society (Demirci et al., 2025; Shan and Qiu, 2025). While the technology holds great potential to enhance productivity and user experiences, it also raises concerns regarding misinformation, authenticity, and consumer protection (Campante et al., 2025). Policymakers around the world have begun to intervene: for example, the Japanese parliament recently passed a law to curb the "rapid spread of false or harmful content generated by AI tools" and granted the government authority to publicly name companies found misusing such technologies.[1]

In parallel, the food delivery industry has grown rapidly over the past decade, reshaping how consumers interact with restaurants (Li and Wang, 2025). Platforms such as Uber Eats, Deliveroo, Foodpanda, Meituan, Zomato, and Swiggy act as intermediaries between restaurants and consumers, enabling convenient online transactions. Uber Eats alone increased its revenue from 1.4 billion USD in 2019 to 12.2 billion USD in 2023.[2] To sustain growth and attract customers in this highly competitive sector, platforms continually experiment with technological innovations to improve user experience.

A recent wave of innovation in the industry comes from the adoption of generative AI. Food delivery platforms are experimenting with AI-powered assistants to guide consumers during the purchase process[3] and with AI-modified dish images that make menus more visually appealing.[4] These applications promise to enhance browsing and shopping experiences, as research shows that images have a significant and important effect in consumer purchase decisions (Li and Xie, 2020). However, they also raise critical questions about consumer trust and perceptions of authenticity. If dish images diverge too far from the food actually delivered, customers may feel misled, avoid reordering from the same restaurant, or even reduce their engagement with the platform. Moreover, concerns about inauthentic images may discourage consumers from trying new restaurants altogether. On Reddit, there have been numerous posts complaining that the dish images "look like AI-generated", and the images shown online differ from the food customers actually receive. Many users also report avoiding restaurants whose dish images appear inauthentic or even boycotting the food delivery platform allowing the usage of generative AI for dish images [5]. Thus, while generative AI offers exciting opportunities for platforms, the impact of having generative AI-manipulated content on consumer behavior and market outcomes remains uncertain. If consumers rely on and value the authenticity of certain content,

[1] See `https://babl.ai/japan-passes-new-ai-law-balancing-innovation-and-risk-mitigation/` (accessed August 2025).

[2] See `https://fourweekmba.com/uber-eats-revenue/` (accessed December 2024).

[3] See `https://tinyurl.com/4cezb3ne` (accessed December 2024).

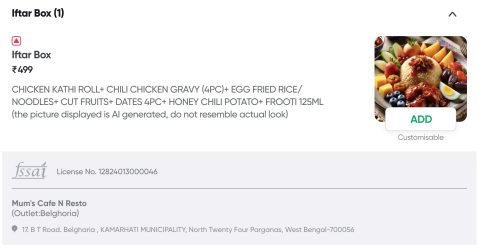
[4] See `https://tinyurl.com/38zsv8tv` (accessed December 2024).

[5] See `https://www.reddit.com/r/mildlyinfuriating/comments/1euhwvl/ai_pics_on_swiggy_are_misleading/` and `https://www.reddit.com/r/swiggy/comments/1nkf3lv/swiggy_should_restrict_restaurants_putting_ai/` (accessed November 2025).

such as dish images, when making purchase decisions, having generative-AI content may even backfire. This makes the study of AI-images valuable and timely, especially in the context where consumers may rely on images for purchase decision-making. Hence, our research question is, *What is the impact of (banning) generative-AI dish images on food delivery platform order volume and revenue? How do consumers react to restaurants having dish images that are more likely to be generative-AI-generated?*

We conduct our empirical analysis using data from India's two largest food-delivery platforms, Zomato and Swiggy, which generated approximately 2 to 3 billion USD in revenue in 2024, respectively.[6] In late 2023, following the global boom of generative AI technologies, both platforms introduced generative AI–based image enhancement tools that allowed restaurant owners to improve dish quality, adjust aesthetics, and choose from a variety of background options.[7] Restaurants could also employ external tools such as Mid-Journey to generate, modify, or blend images.[8] While Swiggy continued permitting restaurants to use AI-generated images throughout 2024 and 2025, Zomato imposed a platform-wide ban on August 18, 2024, citing consumer complaints that AI-modified images were misleading and risked eroding trust. Restaurants that violated the policy faced de-listing from the platform.[9] Figure 1 shows how one Swiggy restaurant utilizes generative AI to display and represent dish images.[10]

Figure 1: One Swiggy Restaurant Disclosing Gen AI Usage



This policy divergence between Zomato and Swiggy offers a unique quasi-experimental setting to study the impact of generative AI on consumer behavior and platform competition. On one hand, AI-enhanced images may attract consumers by improving visual appeal and lowering search costs. On the other hand, if consumers perceive these images as deceptive, they may reduce trust in restaurants or the platform itself. The contrasting strategies of the two platforms allow us to investigate whether the benefits of image enhancement outweigh the potential costs of eroding authenticity, and how such trade-offs shape consumer demand and competitive dynamics in online food delivery markets.

[6] See `https://tinyurl.com/63m7s2ns` and `https://www.jika.io/quote/ZOMATO.NS/financials/revenue` (accessed July 2025).

[7] See `https://tinyurl.com/49mbspbu` and `https://tinyurl.com/2e2atrty` (accessed June 2025).

[8] See `https://tinyurl.com/2tbz5sne` (accessed December 2024).

[9] See `https://tinyurl.com/4yr24ffs` (accessed June 2025).

[10] See `https://www.swiggy.com/city/kolkata/mums-cafe-n-resto-belghoria-rest821126` (accessed October 2025).

Utilizing a unique dataset documenting a random sample of Indian food delivery consumers' purchase records on the two major food-delivery platforms from March 2024 to February 2025, we construct a restaurant-level sales panel. Treating Zomato restaurants as the "treated" group and Swiggy restaurants as the control group, and applying a difference-in-differences approach, we find that restaurants' order volume and revenue on Zomato increased significantly compared with those on Swiggy after the policy change. We also observe that an increasing number of consumers shifted to purchasing from Zomato relative to Swiggy and became more willing to try new restaurants on Zomato. This pattern suggests that eliminating concerns around generative-AI dish images may help restore consumers' confidence in exploring new establishments on the platform.

We further strengthen our analysis by conducting a series of image-forensic tests and multi-modal analyses using state-of-the-art generative-AI image detectors. Forensic examinations of image sharpness, compression-artifact error levels, and GAN fingerprint residuals provide consistent evidence that Swiggy's dish images are more likely to be AI-generated. Three HuggingFace-based generative-AI detectors yield similar results. We also find that Swiggy restaurants with a higher likelihood of using AI-generated images experience larger declines in sales, reflecting consumers' dislike towards images that are more "AI-like" (Exner et al., 2025).

Interestingly, we observe a larger and more significant increase in order volume among non-chain and single-homing restaurants on Zomato, compared with chain and multi-homing restaurants. While platforms may intend for generative-AI tools to help small and independent restaurants—who often lack access to professional photography—consumers may in fact react more negatively to AI-generated images in this segment. One possible explanation is that consumers rely less on dish images when ordering from large, standardized chain restaurants, but place greater weight on image authenticity for small, non-chain restaurants. As a result, AI-generated images may raise concerns about authenticity more acutely for these restaurants and ultimately deter purchases.

This study makes several contributions and has important policy implications. First, we contribute to the literature on human-AI interaction (Castelo et al., 2019; Yang et al., 2025; Krakowski et al., 2025). "Algorithm aversion" theory Castelo et al. (2019) posits that humans would prefer humans over algorithms. For example, Yang et al. (2025) runs a field experiment with a large savings bank in Europe, and finds that bank consumers prefer human-AI joint decisions over pure AI investment advice. We use real-world consumer purchase records and a natural experiment to reveal consumers' preference for real dish images over AI-generated dish images, which can inform platform decision-making and merchants' image selection.

Secondly, we contribute to the growing empirical studies on consumers' perception of generative AI content. Since the launch of ChatGPT in late 2022, generative AI has been transforming the business models of enterprises in many industries. Yet, research exploring consumers' reactions to the content,

especially images, generated by AI in real-world practice has been limited (Belanche et al., 2025; Exner et al., 2025; Huang and Katona, 2025; Li et al., 2025). The closest one is Belanche et al. (2025), which leverages laboratory experiments and found that participants prefer hospitality services advertised with real images over AI-generated images, particularly for hedonic products such as food, compared with utilitarian products. Our empirical evidence from the natural experiment supports their results: food delivery consumers prefer authentic dish images over AI-generated images. Another highly relevant one is Exner et al. (2025), which investigates a quasi-experiment on an advertisement platform and shows that, AI-generated ads receive more clicks compared with human-generated images only if these AI-generated images do not "look like AI". Our empirical findings are consistent with their results: Swiggy restaurants whose dish images have lower AI likelihood experience smaller sales declines, whereas those with higher AI likelihood suffer more substantial reductions in order volume.

Lastly, our research contributes to the growing literature on multi-modal machine learning analysis in the business domain (Chen et al., 2024; Jin et al., 2025; Zhang et al., 2022). While unstructured text data, such as consumer reviews, has been explored in the information systems and marketing literature (Wang et al., 2022), research on images, especially AI-generated images, their alignment with text, and how these affect consumer perception and purchase has been limited. Our research explores several dimensions of dish images based on the computer vision and multi-modal machine learning literature, and analyzes their moderating effect on restaurant sales that future research may find helpful.

2 Dataset and Model-Free Evidence

2.1 Dataset

Our research combines two datasets from different sources. The first is a unique dataset of 13,755 consumers' food delivery purchase records on Zomato and Swiggy in India, obtained from Measurable AI, a leading market intelligence company specializing in tracking consumers' platform transactions. The dataset covers the period from March 1, 2024, to February 28, 2025, and includes purchase records from the same set of consumers across both platforms. It contains rich information such as order value, discounts (Swiggy only), restaurant name, dish purchased, dish price (Swiggy only), and restaurant address. A key feature of this dataset is its ability to observe individual consumers' behavior across platforms, enabling direct cross-platform comparisons. This feature is particularly valuable when implementing our cross-platform analyses. We restricted users to consumers with pre-policy orders to make sure they are treated by the policy change, which leaves us with a final sample of 9723 consumers.[11] Because restaurants on Zomato are the entities

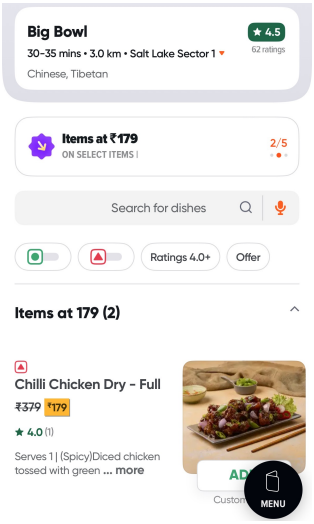
[11] We run a robustness check using data from all the 13,755 consumers and obtained similar results.

directly affected by the policy change, we aggregate consumers' order volumes to obtain restaurant-level sales, which we use as a proxy for restaurant demand.

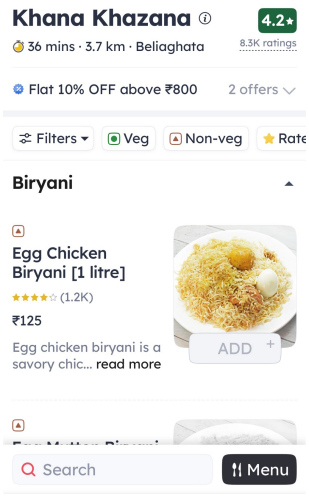
The second dataset consists of restaurant dish images and characteristics that we collected from Zomato and Swiggy after Zomato's generative AI ban in Kolkata, one of the largest cities in India. It contains restaurant information as displayed on the website, including restaurant address, rating, number of reviews, cuisine, price category, dish name, dish price, dish description in English, and dish images. We merged the restaurant sales dataset (restricted to Kolkata) with the restaurant image and characteristics dataset using restaurant name and address when conducting the multi-modal image analysis and mechanism analysis. The sample homepages of two restaurants on the platforms are shown in Figure 2.

Figure 2: Sample Restaurant Pages on Swiggy and Zomato

(a) Swiggy Sample Restaurant Homepage



(b) Zomato Sample Restaurant Homepage



We present summary statistics of restaurants in Table 1. Overall, restaurants on the two platforms are comparable, though Zomato hosts slightly more restaurants, while Swiggy restaurants have marginally higher pre-policy order volumes and lower average order values.[12]

[12]Pre-Policy Average Order Value is missing for restaurants with zero pre-policy orders.

Table 1: **Sample Descriptive Statistics of Restaurants**

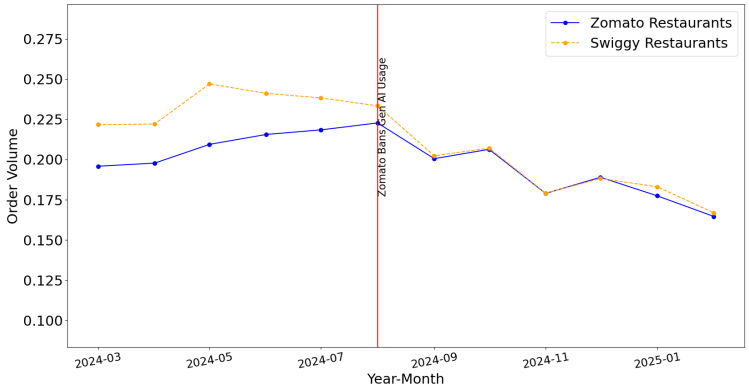
Platform	Variable	Count	Mean	Median	Min	Max
Zomato	Chain (=1)	30,836.00	0.39	0.00	0.00	1.00
Zomato	Pre-Policy Order Volume	30,836.00	1.17	1.00	0.00	49.00
Zomato	Pre-Policy Order Value (INR)	19,068.00	342.86	270.02	19.36	11,039.09
Zomato	Multihoming	30,836.00	0.11	0.00	0.00	1.00
Swiggy	Chain (=1)	27,894.00	0.40	0.00	0.00	1.00
Swiggy	Pre-Policy Order Volume	27,894.00	1.32	1.00	0.00	92.00
Swiggy	Pre-Policy Order Value (INR)	18,343.00	322.83	253.00	0.00	6,080.00
Swiggy	Multihoming	27,894.00	0.12	0.00	0.00	1.00

2.2 Model-Free Evidence

Figure 3 plots average order volume per restaurant on Zomato and Swiggy during the sample period. Before August 2024, which is the month when Zomato banned the use of generative-AI–created dish images (marked by the red line), restaurants' order volume on the two platforms exhibited parallel trends.

Following the ban, however, the trajectories begin to diverge: Zomato restaurants experience a relative increase in order volume compared with restaurants on Swiggy, though Zomato restaurants also have a slight decline due to seasonality as food delivery users naturally buy less after the summer break. This model-free evidence suggests that the ban may have contributed to improved restaurant performance on Zomato, a pattern we examine more formally in the difference-in-differences analysis that follows.

Figure 3: Average order volume per restaurant on two platforms in India



3 Regression Analysis

3.1 Restaurants' Order Volume and Revenue

We conduct a difference-in-difference analysis to assess the impact of the generative AI image ban on restaurants' order volume. The control group is restaurants on Swiggy, while the treatment group is restaurants on Zomato.

The regression model that we use is,

$$\mathrm{Y}_{rt} = \alpha + \beta_1 Treat_r \times PostPolicy_t + \delta_r R_r + \gamma_t T_t + \varepsilon_{rt} \tag{1}$$

where Y_{rt} is restaurants' weekly order volume, total revenue, and number of customers. $Treat_r$ indicates whether a restaurant is a Zomato restaurant or not, $PostPolicy_t$ indicates whether week t is after the policy change or not, R_r is restaurant fixed effect, and T_t is week fixed effect to control seasonality and time trend. Note that for multi-homing restaurants that join both platforms, we treat them as different restaurants on each platform. This is because the same restaurant may offer different dishes, the same dishes at different prices, and, importantly, use different images for the same dish—especially given that they may continue to use generative AI to modify their dish images on Swiggy. The results are displayed in Table 2.

Table 2: **Impact of the Policy Change on Restaurants' Order Volume and Revenue**

	(1) Order Volume	(2) Revenue
Treat × PostPolicy	0.00544***	1.566***
	(0.000802)	(0.302)
Constant	0.0447***	15.41***
	(0.000230)	(0.0867)
Rest FE	Yes	Yes
YearWeek FE	Yes	Yes
Observations	3112690	3112690
R^2	0.086	0.078

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

We find that after Zomato banned restaurants from using generative AI, restaurants on the platform experienced a 12.1% increase in order volume and 10.1% increase in revenue compared to Swiggy. This suggests that consumers likely value authenticity and consistency between dish images and the delivered food more than the visual appeal of the images themselves. These findings indicate that the generative AI ban may enhance Zomato's competitiveness.

3.2 Restaurants' Order Source Analysis

The increase in Zomato restaurants' order volume could stem from two sources. The first is new customers to these restaurants. Previously, with the use of generative AI to modify dish images, customers may have been reluctant to try unfamiliar restaurants because of the risk that the food would look different from the online images. Following the platform-level ban on generative AI for dish images, Zomato consumers may have regained trust in the authenticity of images, making them more willing to explore new restaurants.

The second source is returning customers. In the past, a customer might have been drawn to a restaurant with appealing images but, after receiving food that did not match the photos, chosen not to reorder. With the generative AI ban on dish images, there should be improved consistency between images and actual dishes on Zomato. If customers enjoy the food, they are more likely to reorder from the same restaurant. Nevertheless, this remains an empirical question worth further investigation.

To examine this, we conduct a subgroup analysis categorizing restaurant orders by source—new versus returning customers.[13] To measure traffic to restaurants, we sum the orders from new and returning customers for restaurant r in week t, defining them as $OrderVolumeNewCustomers_{rt}$ and $OrderVolumeReturningCustomers_{rt}$, respectively. These two measures serve as dependent variables in Equation 1. The regression results are presented in Table 3.

Table 3: **Order Source Analysis**

	(1) Order Volume New Customers	(2) Order Volume Returning Customers
Treat × PostPolicy	0.00390***	0.00154**
	(0.000385)	(0.000651)
Constant	0.0243***	0.0204***
	(0.000111)	(0.000187)
Rest FE	Yes	Yes
YearWeek FE	Yes	Yes
Observations	3112690	3112690
R^2	0.013	0.134

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

We find that the policy change has a significant positive impact on Zomato restaurants in attracting both new customers' and returning customers' orders. This suggests that Zomato's generative AI ban boosts consumer traffic by encouraging them to explore new restaurants and also fueling returning purchases. The positive effect is larger and more significant for attracting new consumers' orders, suggesting that, banning

[13] We define a new customer as a first-time buyer to restaurant t; otherwise, they are classified as a returning customer to restaurant t.

AI-generated images platform-wide may help eliminate consumers' concerns on trying new restaurants and encourage them to explore new restaurants and their dishes, potentially improving their utilities.

3.3 Restaurants' Customer Number

We further analyze policy change's impact on consumer traffic. The dependent variables that we employ include number of total consumers, number of new consumers, and number of returning consumers. New and returning consumers are defined as discussed in the previous subsection, depending on whether they are first time buying from restaurant r or not. We re-run regression model 1 and results are presented in Table 4. We find that the policy change has a positive impact on number of total, new, and returning consumers. The effect is larger for the number of new consumers, in line with the order source analysis results.

Table 4: **Consumer Source Analysis**

	Num Total Consumers	Num New Customers	Num Returning Customers
	(1)	(2)	(3)
	Num Customers	Num New Customers	Num Returning Customers
Treat × PostPolicy	0.00507***	0.00390***	0.00117***
	(0.000629)	(0.000385)	(0.000449)
Constant	0.0408***	0.0243***	0.0166***
	(0.000181)	(0.000111)	(0.000129)
Rest FE	Yes	Yes	Yes
YearWeek FE	Yes	Yes	Yes
Observations	3112690	3112690	3112690
R^2	0.069	0.013	0.120

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4 Multi-modal Analysis

4.1 Forensic Tests

To assess whether restaurant dish images were likely generated or enhanced by generative AI, we also conducted a series of forensic tests on a large sample of images collected from Swiggy and Zomato.

We applied three complementary forensic techniques: (i) Error Level Analysis (ELA), which detects inconsistent compression artifacts; (ii) Noise and Frequency Spectrum analysis, which evaluates sharpness and high-frequency texture patterns; and (iii) GAN fingerprint residuals, which identify statistical anomalies characteristic of synthetic image generation. These methods are widely used in digital image forensics and

provide triangulated evidence of AI-based modification (Dzanic et al., 2020; Neves et al., 2022). Figure 4 illustrates these tests for a representative cake image from the same restaurant on Swiggy and Zomato[14]. The Swiggy version shows inconsistent error levels across the cake surface and background props, extremely high sharpness (≈ 225) with dense high-frequency textures, and residual statistics with heavy-tailed distributions (mean = 1.70) and higher ELA (95th percentile is 50). In contrast, the Zomato version exhibits lower sharpness (≈ 113), softer frequency patterns consistent with casual smartphone photography, residuals aligned with typical image compression (mean = 1.29), and lower ELA (95th percentile is 40). Taken together, these differences strongly suggest that the Swiggy image was AI-generated or AI-enhanced (e.g., background replacement, synthetic texture sharpening), while the Zomato image is more likely to be authentic.

Similarly, another group pair is Chicken Biryani provided by a restaurant named Akram's [15], as shown in Figure 5. The swiggy image shows much higher sharpness (1,808 vs. 147), indicating denser high-frequency detail and potential oversharpening from generative or editing tools. They also exhibit higher mean pixel intensity (6.80 vs. 2.60) and greater variation (standard deviation of 9.20 vs. 2.28), suggesting higher dynamic range manipulation and contrast enhancement. The ELA metrics are likewise higher for Swiggy (mean = 41.6, 95th percentile = 150) compared to Zomato (mean = 30.1, 95th percentile = 90), consistent with repeated recompression or partial image synthesis.

Extending this analysis to the full sample of restaurant dish images (Swiggy: 35,083 unique images; Zomato: 51,044 unique images), we find systematic differences across all three forensic measures. On average, Swiggy dish images exhibit substantially higher sharpness (921 versus 622), higher residual statistics (mean of 3.87 versus 3.40), and higher error level analysis (95th percentile of 42.64 versus 34.87). These patterns suggest that Swiggy images are consistently more processed or enhanced relative to their Zomato counterparts.

These results are consistent with our prediction that Swiggy restaurants adopted generative AI-modified images, while Zomato banned them. The discrepancy across platforms is large and robust enough to serve as a useful falsification test in our empirical analysis.

[14]The restaurant name is Cakes, and the dish name is Dry Choco Fudge Cake. Their URLs on the two platforms are `https://www.zomato.com/kolkata/cakes-1-beliaghata/order` and `https://www.swiggy.com/city/kolkata/cakes-beliaghata-rest58090` (accessed September 2025).

[15]Their URLs on the two platforms are `https://www.zomato.com/kolkata/akrams-kestopur/order` and `https://www.swiggy.com/city/kolkata/akrams-kestopur-lake-town-rest779306` (accessed September 2025).

Figure 4: Images of "Dry Choco Fudge Cake" from *Cakes* (Left Swiggy, Right Zomato)




Figure 5: Images of "Chicken Biryani" from *Akram's* (Left Swiggy, Right Zomato)




4.2 Generative AI Image Detection Analysis

To assess whether dish images on Swiggy are more likely to be generated or modified by generative AI relative to Zomato, we applied three independent HuggingFace detection models ("umm-maybe/AI-image-detector," "Smogy/SMOGY-Ai-images-detector," and "Ateeqq/ai-vs-human-image-detector") to all dish images from restaurants on the two platforms. Each detector was applied in batch mode. For each image, we extracted the probability score of being AI-generated ("AI prediction score"). We then averaged these scores across all available dish images within a restaurant, yielding one aggregated score per restaurant per platform. Note that, while the generative AI image detector tends to be reliable, they are not completely accurate (Exner et al., 2025). Our expectation is that, the detectors would predict dish images on Swiggy to be significantly more likely to be AI-generated or AI-modified compared with Zomato, while acknowledging the possible existence of type-I and type-II errors for individual dish images. For example, "umm-maybe/AI-image-detector" claims that its prediction accuracy is around 0.94 on its validation set [16].

We first compared these restaurant-level AI prediction scores between Swiggy and Zomato, separately for two groups of multi-homing restaurants: those displaying *different* images across the two platforms and

[16] See `https://huggingface.co/umm-maybe/AI-image-detector`.

those displaying *identical* images. For restaurants with different images, at the dish image level, of all the 3153 Swiggy dish images, $median_{swiggy} = 0.33$, and of all the 2342 Zomato dish images, $median_{zomato} = 0.04$. At the restaurant level, the average AI prediction score is significantly higher for Swiggy compared to Zomato ($N = 34$, $t = 5.64$, $p < 0.001$, $median_{swiggy} = 0.32$, $median_{zomato} = 0.15$). This result is consistent with our prediction that Swiggy is more likely to adopt AI-modified or AI-generated dish images. In contrast, for restaurants that display identical images across the two platforms, the difference is statistically insignificant ($N = 144$, $t = -0.95$, $p = 0.34$). For single-homing restaurants, we find similar results ($N_{Swiggy} = 478$,$N_{Zomato} = 797$, $t = 10.25$, $p < 0.001$).[17]

Taken together, these results provide evidence that the observed discrepancies in dish images between Swiggy and Zomato are systematically related to the use of generative AI, rather than random variation in photography or presentation style. In particular, when a restaurant reuses the same image across platforms, there is no systematic difference in AI likelihood, but when images differ, Swiggy's version is much more likely to carry generative AI signatures.

5 Mechanism Analysis

5.1 Multi-homing and Single-homing Restaurants

Theoretically, multi-homing restaurants, which are restaurants available on both platforms, have less incentive to use generative AI images on Swiggy for consistency reasons. In addition, consumers may also verify dish images across platforms for multi-homing restaurants to check the authenticity of dish images. Therefore, even if they intended to beautify dish images on Swiggy, the effect would be limited or null. As a result, we predict that policy change should not have a significant impact on their order volume. By contrast, single-homing restaurants, which are restaurants that are available on only one platform, are more likely to employ generative AI images on Swiggy as they do not have consistency concerns, and consumers are not able to verify dish images across platforms. These restaurants are likely to be more affected by the policy change. We re-run regression model 1 based on restaurants' multi-homing status, and the results are displayed in Table 5, which are in line with our predictions.

[17]Restaurants with missing AI scores on Swiggy are excluded from the analysis. Average AI score is missing for restaurants without dish images and restaurants whose URLs are not available. Among all the 1277 Zomato restaurants, 247 of their URLs are no longer available. Among all the 1251 Swiggy restaurants, 494 of their URLs are no longer available. The average AI score for each restaurant is computed from three independent HuggingFace detectors.

Table 5: **Order Volume Analysis by Restaurants' Multi-homing Status**

	Single-Homing Restaurants	Multi-Homing Restaurants
	(1)	(2)
	Order Volume	Order Volume
Treat × PostPolicy	0.00567***	0.00367
	(0.000832)	(0.00275)
Constant	0.0424***	0.0616***
	(0.000239)	(0.000778)
Rest FE	Yes	Yes
YearWeek FE	Yes	Yes
Observations	2749004	363686
R^2	0.082	0.103

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The results for single-homing restaurants remain positive and significant. This also alleviates concerns regarding the Stable Unit Treatment Value Assumption (SUTVA). For multi-homing restaurants, sales performance across the two platforms is highly correlated—restaurants that perform well on one platform tend to perform well on the other. Their outcomes may therefore be jointly driven by unobservable restaurant-level demand shocks rather than by the platform's image policy. The consistent results between single-homing and the overall sample show the robustness of our results.

5.2 Alternative Control Group Categorized by Dish Image AI Likelihood

To further examine whether the post-policy effects differ depending on the extent of AI involvement in restaurants' imagery, we classify Swiggy restaurants into *low-* and *high-AI-likelihood* groups using the median Swiggy restaurant-level AI prediction score (0.34) as the cutoff. All Zomato restaurants remain in both regressions as the treated group, while the composition of the control group varies depending on whether a Swiggy restaurant's average AI score falls below or above the cutoff. We then estimate Equation 1 separately for each subsample. We restrain the analysis into Kolkata sample for this and following subsections, which is the city where we have restaurant dish image dataset.

A simple subgroup split—pooling Zomato and Swiggy restaurants and classifying both using a single AI-likelihood cutoff—would generate an unbalanced comparison. Because Swiggy restaurants exhibit systematically higher AI manipulation likelihood than Zomato restaurants (which were constrained by the ban beginning August 2024), a naive cutoff would place disproportionately more Swiggy restaurants into the high-AI-likelihood group and disproportionately more Zomato restaurants into the low-AI-likelihood group. This imbalance would confound the subgroup comparison and bias the estimation. Our design instead varies

the *control* group composition while holding the treated group fixed, allowing us to isolate how the treatment effect differs depending on the AI intensity of Swiggy restaurants' imagery. Conceptually, this setup compares Zomato restaurants—none of which can use generative AI after the ban—to more or less "AI-looking" restaurants on Swiggy.

Table 6 presents the results. When Swiggy restaurants with low AI-likelihood images are used as the control group (Column 1), the DiD coefficient on the interaction term *Treat* × *PostPolicy* is smaller in magnitude and less statistically significant ($\beta = 0.0182$, $p < 0.05$). This suggests that, relative to Swiggy restaurants whose images appear largely authentic, Zomato restaurants experience a modest post-policy lift in order volume.

In contrast, when the control group consists of Swiggy restaurants with high-AI-likelihood images (Column 2), the DiD estimate becomes economically larger and statistically stronger ($\beta = 0.0239$, $p < 0.01$). This indicates that restaurants on Swiggy that rely more heavily on AI-enhanced imagery exhibit a more pronounced decline in order volume following Zomato's generative-AI ban.

Together, these findings reinforce the interpretation that the positive post-policy impact for Zomato is concentrated among market segments where AI-generated or AI-modified imagery is more prevalent. The results are consistent with the notion that consumers respond more negatively to dishes that appear "more AI-like" once concerns about image authenticity become salient, echoing behavioral evidence that consumers tend to distrust algorithmically generated content when authenticity is valued (Dietvorst et al., 2015). In other words, Zomato's ban appears to have strengthened consumer confidence in the platform precisely where AI manipulation was most pervasive, thereby amplifying the competitive advantage against high-AI-likelihood restaurants on Swiggy.

Table 6: **Moderating Effect of Dish Image AI Likelihood**

	Low-AI-Likelihood Restaurants as Control	High-AI-Likelihood Restaurants Control
	(1)	(2)
	Order Volume	Order Volume
Treat × PostPolicy	0.0182**	0.0239***
	(0.00922)	(0.00811)
Constant	0.0386***	0.0362***
	(0.00421)	(0.00372)
Restaurant FE	Yes	Yes
YearWeek FE	Yes	Yes
Observations	63971	63759
R^2	0.065	0.068

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

5.3 Alternative Control Group: Multi-homing Restaurants with Different Dish Images

One potential concern with our baseline results is that factors unrelated to dish images—such as platform-specific promotions, algorithmic rankings, or seasonal customer behavior—may partly explain the observed divergence in post-policy order volume across platforms. Because we cannot directly observe whether a restaurant actually used Generative AI to modify its dish images, we exploit cross-platform variation among multi-homing restaurants to build an alternative, more conservative control group.

Specifically, we identify Swiggy restaurants in Kolkata that list on both platforms but display different dish images from those shown on their corresponding Zomato listings. Theoretically, restaurants operating on multiple platforms face strong incentives to maintain consistent visual branding: using the same image set minimizes photography and editing costs, ensures brand coherence, and avoids confusing consumers. Therefore, when a single restaurant exhibits divergent images on Swiggy relative to Zomato after Zomato's generative-AI ban, it is plausible that the Swiggy images were modified with Generative AI tools and thus prohibited on Zomato. This creates a natural, behaviorally grounded contrast between AI-constrained listings (Zomato) and potentially AI-enhanced listings (Swiggy).

We identify 21 such multi-homing Swiggy restaurants in Kolkata whose dish images differ from those displayed by the same entity (matched by restaurant name and address) on Zomato. These 21 Swiggy restaurants serve as the alternative control group, while all Zomato restaurants in Kolkata remain the treated group. We re-estimate Equation 1 for the three main dependent variables—order volume, total revenue, and number of customers.

Importantly, we do not restrict the treated group to only multi-homing Zomato restaurants. Doing so would risk violating the Stable Unit Treatment Value Assumption (SUTVA) because a multi-homing restaurant's sales performance on the two platforms is highly correlated. Restaurants that are popular on one platform tend also to be popular on the other, and their cross-platform outcomes may be jointly determined by unobservable restaurant-level demand shocks rather than the platform's image policy. Using such restaurants simultaneously as "treated" and "untreated" units would therefore confound the identification strategy and bias the DiD comparison. Retaining the full set of Zomato restaurants as the treated group avoids these issues.

The resulting estimates, summarized in Table 7, consistently show positive and statistically significant post-policy effects across all three outcome measures. These results reinforce our main conclusion: the policy banning AI-generated dish images produced meaningful increases in consumer demand for Zomato restaurants, even when benchmarked against a narrow, plausibly AI-intensive comparison group.

Table 7: **Alternative Control Group - Swiggy Restaurants that Have Different Dish Images**

	(1) Order Volume	(2) Revenue	(3) Num Customers
Treat × PostPolicy	0.0511***	17.55***	0.0368**
	(0.0194)	(6.566)	(0.0160)
Constant	0.0190*	9.491***	0.0231***
	(0.0104)	(3.520)	(0.00857)
Restaurant FE	Yes	Yes	Yes
YearWeek FE	Yes	Yes	Yes
Observations	54537	54537	54537
R^2	0.063	0.076	0.054

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

5.4 Alternative Control Group: Multi-homing Restaurants that have Same Dish Images

As a complementary robustness check, we examine a second group of multi-homing restaurants—those that display identical dish images on both Swiggy and Zomato. Theoretically, when a restaurant uses the same image set across platforms, its exposure to Zomato's generative-AI ban should be minimal. If the images remain unchanged after the policy, the restaurant's order volume should not systematically respond to the ban through the image-authenticity mechanism. In other words, this subgroup should behave like an "untreated" control group with respect to the image policy, providing a natural placebo test.

Using our cross-platform image-matching procedure, we identify 84 Swiggy multi-homing restaurants that display identical dish images on both platforms. Because these restaurants do not appear to rely on platform-specific image manipulation, comparing their Swiggy performance with all Zomato restaurants allows us to test whether any residual platform-level differences (unrelated to imagery) are driving the main results. If our baseline findings are indeed driven by consumers' responses to image authenticity rather than by unrelated confounders, the DiD estimates using this "same-image" control group should be statistically indistinguishable from zero.

We re-estimate Equation 1 for all three outcome variables—order volume, total revenue, and number of customers—using this set of 84 restaurants as the Swiggy control group. The results, shown in Table 8, confirm this interpretation: the interaction term *Treat × PostPolicy* is small in magnitude and statistically insignificant across all outcomes. This pattern provides additional validation that our main findings are not driven by platform-wide shocks, differential promotions, or sample composition differences. Rather, the

effects emerge precisely when Swiggy restaurants present imagery that is plausibly AI-modified or visually distinct from the constrained images allowed on Zomato.

Table 8: **Alternative Control Group - Swiggy Restaurants that Have the Same Dish Images**

	(1) Order Volume	(2) Revenue	(3) Num Customers
Treat × PostPolicy	-0.000818	-1.265	-0.00242
	(0.0101)	(5.575)	(0.00878)
Constant	0.0477***	20.21***	0.0450***
	(0.00509)	(2.816)	(0.00443)
Restaurant FE	Yes	Yes	Yes
YearWeek FE	Yes	Yes	Yes
Observations	57876	57876	57876
R^2	0.065	0.080	0.057

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

5.5 Correlation between Swiggy Restaurant Ratings and AI Likelihood

Anecdotal observations suggest that consumers tend to respond negatively to restaurants whose dish images appear to be AI-generated [18]. To provide quantitative support for this pattern, we conduct a correlation analysis between the AI likelihood score of restaurant images on Swiggy and their corresponding average customer ratings. After removing missing values, we find a statistically significant negative association. The Pearson correlation between the AI likelihood score and Swiggy ratings is $-0.068(N = 661, p = 0.079)$, and the Spearman rank correlation is $-0.091(N = 661, p = 0.019)$[19]. These results indicate that restaurants with higher probabilities of using AI-generated images tend to receive slightly lower customer ratings, consistent with the notion that consumers may perceive such images as less authentic and unfavorable.

6 Heterogeneous Analysis of Chain and Non-Chain Restaurants

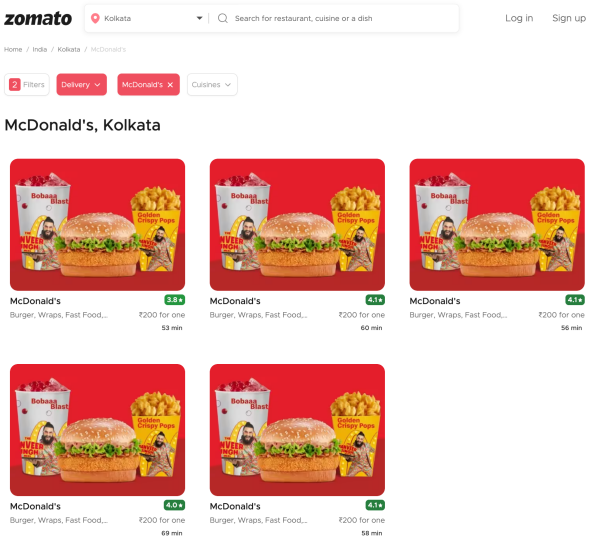
Theoretically, while small and independent restaurants have more freedom and flexibility to utilize generative AI technology, chain restaurants have less flexibility when it comes to altering or improving their images. This is because dish images for chain restaurants are typically designed and standardized by headquarters. Moreover, these chain restaurants may be more cautious about employing generative AI to enhance their

[18] See `https://www.reddit.com/r/swiggy/comments/1nkf3lv/swiggy_should_restrict_restaurants_putting_ai/` (accessed November 2025).

[19] Spearman rank correlation is more robust to outliers (Sedgwick, 2014).

visuals. For instance, all McDonald's outlets on Zomato in Kolkata display the same header and dish images, as shown in Figure 6. Similarly, the fast-food behemoth uses identical header and dish images across its branches on Swiggy. This consistency suggests that McDonald's likely does not rely on generative AI for its food photography.

Figure 6: Sample Pages of McDonald's in Kolkata on Zomato



In addition to the low likelihood of chain restaurants using generative AI, customers are usually more familiar with the food offered by chain restaurants. Even if these restaurants do use generative AI, the impact on customer purchases is likely limited, as customers already have clear expectations about the food they will receive when ordering online. Therefore, we predict that the policy change would not significantly affect chain restaurants' order volumes, but would primarily impact non-chain restaurants. To test this, we run a separate regression using Equation 1, splitting the sample into chain and non-chain restaurant subgroups. The results are displayed in Table 9.

We find that the regression results align with our prediction. The policy change has a smaller positive impact on chain restaurants, while the number of orders from non-chain restaurants has increased at a larger magnitude and more significantly. These results have two implications. First, they empirically show that chain restaurants are less likely to adopt and benefit less from the use of generative AI technology for image modification. Second, banning the use of generative AI technology improved non-chain restaurants' order volume and market share. This result may seem counterintuitive, as small restaurants often lack the resources and access to professional photographers to take high-quality dish photos, and generative AI technology could ostensibly help them. However, it appears that consumers value authenticity more, and banning the use of generative AI ultimately across the platform benefits non-chain restaurants and improves their competitiveness.

Table 9: **Analysis on Restaurants' Chain Status**

	Chain Restaurants	Non-Chain Restaurants
	(1) Order Volume	(2) Order Volume
Treat × PostPolicy	0.00394***	0.00643***
	(0.00134)	(0.000997)
Constant	0.0519***	0.0399***
	(0.000383)	(0.000287)
Rest FE	Yes	Yes
YearWeek FE	Yes	Yes
Observations	1231508	1881182
R^2	0.093	0.080

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7 Robustness Checks and Alternative Explanations

7.1 Parallel Trend Assumption Check

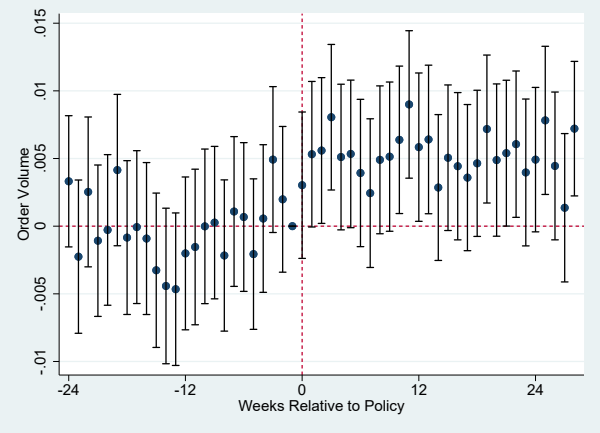
The difference-in-difference analyses rely on the parallel trend assumption. To check the validity of the assumption, following Aral and Dhillon (2021), we create a series of week dummies to indicate the weeks before and after the policy change during the sample period. We then interact each of the dummy variables with the treatment variable and run the following regression:

$$\mathrm{Y}_{rt} = \alpha + \sum_{\tau=-24,\ldots,-2}^{0,1,\ldots,28} \beta_\tau 1[t - event_t = \tau] \times Treat_r + \delta_r R_r + \gamma_t T_t + \varepsilon_{jt} \tag{2}$$

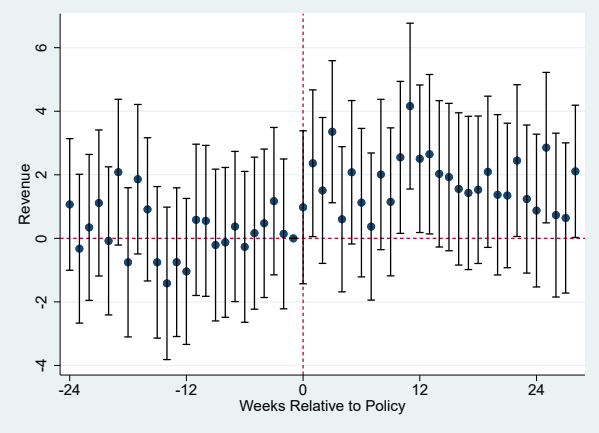
where Y_{rt} includes restaurants' weekly order volume and total revenue. $event_t$ represents the week in which the policy shock occurred. $1[t - event_t = \tau]$ are a series of "event-time" dummies that equal one when the policy shock is τ periods away from event t. The week before the policy shock serves as the baseline, and its event-time dummy variable is omitted. We run the regression for the whole sample in India and restaurant sample in Kolkata seperately. The regression results are shown in 7. As shown in the regression results, the coefficients are not significantly different from 0 before the policy change, indicating that Swiggy and Zomato restaurants do not have a significant difference between order volume and revenue, lending support for our difference-in-difference analysis.

Figure 7: Parallel Trend Assumption Check Results

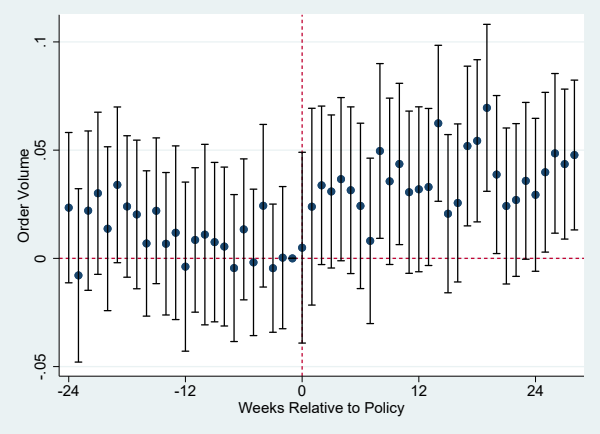
(a) Order Volume - India



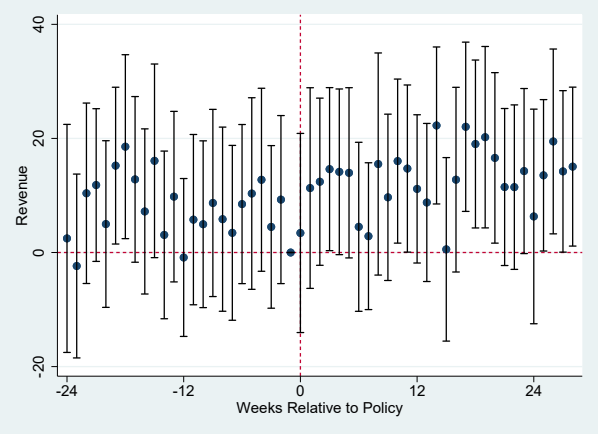
(b) Revenue - India



(c) Order Volume - Kolkata



(d) Revenue - Kolkata



7.2 Promotions and Dish Price on Swiggy

To ensure that the observed changes in platform choice are not confounded by concurrent price or promotion adjustments, we further conduct a robustness check examining whether Swiggy's average order discounts and average dish prices exhibit any discontinuous changes around the policy implementation week.[20] Specifically, we estimate the Regression Discontinuity Design (RDD) model:

$$Y_{rt} = \beta_0 + \beta_1 \text{Weeks to Policy Change}_t + \beta_2 \text{PostPolicy}_t + \beta_3 (\text{Weeks to Policy Change}_t \times \text{PostPolicy}_t) + \alpha_r + \varepsilon_{rt}, \tag{3}$$

where Y_{rt} denotes either the average order discount amount (in Indian Rubee) or average dish price for restaurant r in week t, Weeks to Policy Chang$_t$ is the running variable measuring weeks relative to the policy cutoff, and α_r captures restaurant fixed effects. The coefficient β_2 represents the discontinuity at the

[20] Given we do not have promotion information on Zomato, we cannot run a DiD analysis on these two variables.

cutoff, indicating whether Swiggy changed its discount or pricing policy immediately after the policy took effect.

Table 10 reports the results. The estimated discontinuities at the cutoff are small and statistically insignificant for both outcomes: for the average discount, $p = 0.408$, while for the average dish price, $p = 0.164$. Neither effect is statistically significant at conventional levels. These results suggest that Swiggy did not systematically adjust its discounting or pricing strategies around the policy implementation period.

Taken together, this robustness check supports our identification strategy by confirming that the observed shifts in consumer platform choices are unlikely to be driven by concurrent changes in Swiggy's pricing or discount policies, but rather by the Zomato generative AI ban policy intervention itself.

Table 10: **Robustness Check on Avg Order Discounts and Avg Dish Price**

	(1) Avg Discount Amount Swiggy	(2) Avg Dish Price Swiggy
Months to Policy Change	0.0196	-3.901
	(2.003)	(3.159)
PostPolicy	1.270	8.716
	(4.745)	(7.053)
PostPolicy × Months to Policy Change	-2.520	4.932
	(2.606)	(3.655)
Constant	112.6***	197.4***
	(4.037)	(6.257)
Restaurant FE	Yes	Yes
Observations	2402	3605
R^2	0.769	0.749

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7.3 Order Size

We perform an additional robustness analysis to rule out the possibility that the observed increase in order volume is driven by temporary promotions or platform- or restaurant-level pricing changes. Specifically, we re-estimate Equation 1 using the average number of dishes per order as the dependent variable. If the post-policy increase in order volume were primarily due to promotions (e.g., "1-for-1" deals or discounts), we would expect to observe more dishes per order on Zomato.

As shown in Table 11, the estimated treatment effect on the number of dishes per order is statistically insignificant. These results suggest that the policy change did not affect order size, reinforcing that the observed rise in order volume on Zomato is unlikely to be driven by promotional activities.

Table 11: **Impact of the Policy Change on Rests' Order Size**

	(1) Avg Number of Dishes Per Order
Treat × PostPolicy	-0.0461
	(0.0321)
Constant	2.365***
	(0.00881)
Restaurant FE	Yes
YearWeek FE	Yes
Observations	93482
R^2	0.538

Notes: Robust standard errors, clustered at the restaurant level, are presented in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7.4 User-Level Regression

So far, we have been presenting restaurant-level regression results to demonstrate the policy change's impact on restaurant sales. To check if a user-level regression would yield similar results, we implement a simple reduced-form model:

$$Y_{it} = PostPolicy_t + U_i + \varepsilon_{it} \tag{4}$$

where Y_{it} includes user i's monthly number of order difference between Zomato and Swiggy (number of Zomato orders minus number of Swiggy orders), which measures the order volume change between platforms, U_i is user fixed effect, and T_t is year-month fixed effect. The results are displayed in Table 12. Consistent with previous restaurant-sales results, at the platform level, we find that food delivery users shifted their orders from Swiggy to Zomato.

Table 12: **Impact of the Policy Change on Users' Consumption on Two Platforms**

	Order Volume Zomato minus Order Volume Swiggy
PostPolicy	0.0648***
	(0.0197)
Constant	-0.0148
	(0.0115)
User FE	Yes
Observations	115848
R^2	0.623

Notes. Robust standard errors are clustered at the user level and reported in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

8 Discussion

Leveraging a natural experiment and difference-in-difference methods, we find that banning AI-generated dish images would increase food delivery platform order volume significantly, compared with the platform allowing the usage of them. In addition, we find that compared with Swiggy restaurants with a higher AI-likelihood of dish images, Zomato restaurants tend to have a larger order volume drop. Our research also finds that independent restaurants may benefit more from the generative-AI image ban, as consumers may rely more heavily on image authenticity when buying from them, compared with chain and standardized restaurants.

This research has important theoretical contributions. First of all, the paper adds to existing literature discussing the trade-off between algorithm usage and consumer trust. There has been extensive literature discussing the usage of algorithms and consumers' perceptions, for example, Kim and Duhachek (2020). We contribute to the literature by analyzing the new technology advancement of generative AI on food delivery consumer perceptions. Additionally, we also contribute to the literature on platform competition from the perspective of technology usage (Armstrong, 2006). With the advancement of technologies in the field of generative AI, traditional platform competition literature may need consider how technology innovation and platform adoption may affect the landscape of platform competition. This project explores how the ban on generative AI affects platform competition and marketplace dynamics. Moreover, our research contributes to business studies in emerging markets, particularly South Asia. As one of the largest markets in the world, India and similar economies remain underexplored in terms of empirical research on business models and consumer behavior due to data limitations. Such research is critically important given the region's growing influence in the global market.

Our research also has important managerial implications. Practically, business establishments such as restaurants may be incentivized to use generative AI tools for dish images because these tools can save time and produce more visually appealing photos. However, in contexts where consumers rely heavily on the authenticity of images to inform their decisions, the use of AI-generated visuals may backfire and discourage purchases. From the platform's perspective, widespread use of generative-AI dish images can dilute the platform's brand and undermine consumer trust. Moreover, even restaurants with low AI-likelihood images may suffer if consumers perceive the platform as generally lacking authenticity, potentially driving users toward competing platforms that offer more trustworthy visual content. More broadly, businesses should be cautious when deploying AI-generated content in settings where consumers place a high value on authenticity.

This project is not without limitations. One key limitation is that we do not observe the ground truth which restaurants used generative AI to modify dish images prior to the policy change. While this limitation may not affect the validity of our cross-platform analysis, future research with more detailed datasets could

explore potential spillover effects in order volume changes between restaurants that adopted the technology and those that did not after a generative AI ban.

Our paper opens several avenues for future work. Firstly, while both this research and Exner et al. (2025) find that consumers react more negatively to images that appear more "AI-like," there may be more nuanced heterogeneity across different degrees of AI manipulation, such as AI-slightly-modified and completely-AI-generated. Future work with access to field-experiment settings could collaborate with platforms to examine these gradients more precisely. Second, while prior studies have explored consumer perceptions of AI-generated images in advertising and food-delivery contexts, research in other domains remains relatively sparse. Future research could investigate whether similar mechanisms operate in industries where authenticity is also valued, such as hotels and short-term rentals. Lastly, our research focuses on the impact of AI-images on restaurant orders. Future research may explore other AI-generated content, such as text and videos and check whether similar conclusions would yield for other modalities of AI-generated content.

Appendix

References

Aral, S. and Dhillon, P. S. (2021). Digital paywall design: Implications for content demand and subscriptions. *Management Science*, 67(4):2381–2402.

Armstrong, M. (2006). Competition in two-sided markets. *The RAND journal of economics*, 37(3):668–691.

Belanche, D., Ibáñez-Sánchez, S., Jordán, P., and Matas, S. (2025). Customer reactions to generative ai vs. real images in high-involvement and hedonic services. *International Journal of Information Management*, 85:102954.

Campante, F. R., Durante, R., Hagemeister, F., and Sen, A. (2025). Genai misinformation, trust, and news consumption: Evidence from a field experiment. Technical report, National Bureau of Economic Research.

Castelo, N., Bos, M. W., and Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of marketing research*, 56(5):809–825.

Chen, G., Huang, L., Xiao, S., Zhang, C., and Zhao, H. (2024). Attending to customer attention: A novel deep learning method for leveraging multimodal online reviews to enhance sales prediction. *Information Systems Research*, 35(2):829–849.

Demirci, O., Hannane, J., and Zhu, X. (2025). Who is ai replacing? the impact of generative ai on online freelancing platforms. *Management Science*.

Dietvorst, B. J., Simmons, J. P., and Massey, C. (2015). Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology: General*, 144(1):114.

Dzanic, T., Shah, K., and Witherden, F. (2020). Fourier spectrum discrepancies in deep network generated images. *Advances in neural information processing systems*, 33:3022–3032.

Exner, Y., Hartmann, J., Netzer, O., and Zhang, S. (2025). Ai in disguise-how ai-generated ads' visual cues shape consumer perception and performance. *Available at SSRN*.

Huang, M. and Katona, Z. (2025). Consumer evaluation and visual attribution of realistic ai-generated advertisements. *Available at SSRN 5366745*.

Jin, Y., Chakraborty, I., and Chung, K. (2025). Advancing marketing video analysis with multimodal llms and context-aware graphrag. *Available at SSRN 5517439*.

Kim, T. W. and Duhachek, A. (2020). Artificial intelligence and persuasion: A construal-level account. *Psychological science*, 31(4):363–380.

Krakowski, S., Haftor, D., Luger, J., Pashkevich, N., and Raisch, S. (2025). Human-centered artificial intelligence: A field experiment. *Management Science*.

Li, Y. and Xie, Y. (2020). Is a picture worth a thousand words? an empirical study of image content and social media engagement. *Journal of marketing research*, 57(1):1–19.

Li, Z., Nie, J., Zhang, X., Shi, Z. M., and Li, K. (2025). How do ai-generated summaries affect consumer decision-making? field and experimental evidence from digital deal platforms. *Field and Experimental Evidence from Digital Deal Platforms (July 18, 2025)*.

Li, Z. and Wang, G. (2025). On-demand delivery platforms and restaurant sales. *Management Science*, 71(7):5788–5804.

Neves, J. C., Tolosana, R., Vera-Rodriguez, R., Lopes, V., Proença, H., and Fierrez, J. (2022). Gan fingerprints in face image synthesis. In *Multimedia Forensics*, pages 175–204. Springer Singapore Singapore.

Sedgwick, P. (2014). Spearman's rank correlation coefficient. *Bmj*, 349.

Shan, G. and Qiu, L. (2025). Examining the impact of generative ai on users' voluntary knowledge contribution: Evidence from a natural experiment on stack overflow. *Information Systems Research*.

Wang, H., Du, R., Shen, W., Qiu, L., and Fan, W. (2022). Product reviews: A benefit, a burden, or a trifle? how seller reputation affects the role of product reviews. *MIS quarterly*, 46(2):1243–1272.

Yang, C., Bauer, K., Li, X., and Hinz, O. (2025). My advisor, her ai, and me: Evidence from a field experiment on human–ai collaboration and investment decisions. *Management Science*.

Zhang, S., Lee, D., Singh, P. V., and Srinivasan, K. (2022). What makes a good image? airbnb demand analytics leveraging interpretable image features. *Management Science*, 68(8):5644–5666.