

Will it Click? Advertising Strategy with Generative AI*

Itay P. Fainmesser[†]

Dominique Lauga[‡]

Elie Ofek[§]

This Draft: February 28, 2026.

Abstract

Increasingly, marketers are leveraging AI technologies to interact with consumers. In particular, many companies have adopted recent advent of commercially available Generative AI (GenAI) capabilities (e.g., Open AI's ChatGPT and Google's Gemini), which allow creating personalized messages at scale, to more effectively advertise their offerings. In this paper, we model the impact of GenAI on advertising strategy. Firms seek to sell a product to consumers with heterogeneous tastes and can imperfectly identify their message preferences, i.e., which creative is most likely to prompt a response (or click). The firm has a baseline advertising message template and can use GenAI to adapt it and tailor a message to each consumer. However, the GenAI messages produced can deviate from the intended consumer's preference. In a monopoly setting, we find that the optimal advertising strategy will result in a click probability exhibiting an inverted U shape, depending on the consumer's message preference location along a Salop circle. We also uncover an interplay between the ability to identify consumers' message preference locations and the GenAI technology. Specifically, if the GenAI technology is not too accurate, the optimal size of the segment the firm targets can be initially decreasing (i.e., advertising to fewer consumers) as the ability to identify consumers improves. The model is expanded to analyze the interaction between two competing firms armed with GenAI capabilities and baseline advertising messages located at opposite ends of the message preferences space. We find that if the consumer is located closer to (farther from) a firm in the message space, counterintuitively, the probability of clicking on a firm's ad will decrease (increase) as GenAI's accuracy improves. As a result, better GenAI capability intensifies competition and leads to reduced profits; while the reverse is true as the ability to identify consumer preferences improves. We further endogenize GenAI capability and find that firms' investments in this technology form strategic substitutes, giving rise to an asymmetric equilibrium in which one firm invests more in GenAI capability than its rival. Lastly, when a firm can also identify heterogeneous consumer preferences in the product space, prices tend to rise, as the firm can concentrate on higher willingness to pay consumers. However, for a given level of product preference identification, improvements in GenAI accuracy may prompt a firm to reduce price, as it becomes advantageous to increase advertising reach and thus try to convert more price sensitive consumers. The findings have important managerial implications for firms as they seek to benefit from GenAI technologies, and provide guidance on how to formulate advertising strategies when employing them.

*The authors wish to thank participants at the Zero Decade Conference (London Business School, July 2025) and the University of Maryland Marketing Camp (September 2025) for their valuable comments and suggestions.

[†]The Johns Hopkins Carey Business School and The Department of Economics, The Johns Hopkins University, Baltimore, MD 21202 (e-mail: itay_fainmesser@jhu.edu)

[‡]Cambridge Judge School of Business, University of Cambridge, Cambridge, UK (e-mail: d.lauga@jbs.cam.ac.uk)

[§]Harvard Business School, Harvard University, Boston, MA 02163 (e-mail: eofek@hbs.edu)

1. Introduction

Recent advances in Artificial Intelligence (AI) have taken the business world by storm, with the vast majority of corporations today (over 80%) either actively implementing, or seriously contemplating integrating, AI technologies into their core business functions (Gil and Raymond, 2025). More specifically, the adoption of Generative AI (GenAI) systems, which can create novel outputs — such as text, images, or video — by learning patterns from large unstructured datasets, is gaining commercial momentum at a breakneck pace. GenAI allows firms to automate and augment creative and cognitive tasks at a scale never imagined before. Not surprisingly, therefore, since the introduction of OpenAI’s ChatGPT a little over three years ago, the global GenAI market has already been valued at \$62.75 billion and is expected to grow to \$356 billion by 2030 (a CAGR of 41.5%; Ngugi, 2025)

Several business use cases for GenAI have shown promise, ranging from finance to supply chain management to human resource planning. Among the most highly touted applications is in the domain of marketing communications. Notably, by analyzing individual user behavior and characteristics, and combining that with vast amounts of training data to learn content patterns, GenAI is increasingly being applied to create ads that are targeted and highly personalized. Such personalized ads have been extolled by industry observers as more likely to grab consumers’ attention and result in clicks (Elias, 2025). In fact, McKinsey has recently highlighted examples of how companies, e.g., in the telecom sector, were able to boost response rates to offers by using GenAI tools to customize the ad copy for each individual (Stein et al., 2025). Their observations detailed how pioneering firms first use AI to determine for each individual a “content propensity” measure, which predicts the type of content and the manner in which it is presented that would best encourage an action (such as clicking on an ad). These firms then leverage GenAI to blend general company messaging with individualized modifications informed by the content propensity measure, thereby tailoring the message to each recipient. In another noteworthy instance, financial services firm Vanguard reportedly utilized a GenAI system to produce and deploy personalized ads on LinkedIn, achieving a 15% increase in click through rates (Grewal et al., 2025).

In addition to these practitioner accounts, a number of academic studies have shown in controlled settings the impact GenAI can have on ad effectiveness. For example, Kapoor and Kumar (2025) ran a field experiment in partnership with an e-commerce brand that sells sustainable products and found that personalized video ads produced with the help of GenAI outperformed a non-personalized

generic version by 9%. And a behavioral study by [Matz et al. \(2024\)](#) demonstrated that messages promoting products such as iPhones, which were crafted by ChatGPT and customized to people’s psychological profiles (based, among other things, on their browsing behavior, social media posts, and smartphone use), were rated as more persuasive than non-GenAI created messages.

Despite the immense opportunities AI presents for marketers, in both helping identify which content elements and structures might resonate with each consumer, as well as leveraging GenAI to produce customized ads based on these preferences, and despite a host of encouraging success stories, many companies face challenges in embracing AI. For instance, upwards of 40% of businesses report concerns over using GenAI in consumer facing applications, such as targeted ads, for fear of inaccuracies and hallucinations. Moreover, a significant portion of marketing professionals who have adopted AI tools admit that they struggle to maximize its value and justify the investment ([Ngugi, 2025](#)).

Given these trends in the marketplace, in this paper we aim to shed light on the implications of leveraging AI capabilities, which allow identifying consumer preferences as well as generating customized messages, for advertising purposes. In particular, we seek to understand the forces governing firms’ strategic behavior in the presence of this transformative technology by exploring the following research questions.

- How does the firm’s optimal advertising strategy, when utilizing AI, affect consumers’ probability of clicking on an ad? And how is this probability affected by AI’s accuracy in identifying consumer message preferences and GenAI’s ability to accurately personalize ads?
- How does the size of the market the firm targets depend on the effectiveness of different AI capabilities (message preference identification and content generation)? Should a firm always advertise to more consumers when these capabilities improve?
- In a competitive context, when firms can leverage AI to identify consumer message preferences and personalize ads:
 - How do equilibrium outcomes and profits depend on the capabilities of the prevailing AI technology?
 - When investment in the GenAI capability is endogenous, will competing firms always invest the same amount, or might they invest asymmetrically?
- If the firm can also identify consumers’ product preferences, how does this affect its decision of whom to target with ads? What is the relationship between the firm’s ad-related AI capabilities and the price it charges?

To address these research questions, we construct a model in which a firm seeks to sell its product to consumers. However, in order to become aware of the product and consider a purchase, a consumer must first encounter an ad, which is costly for the firm to send, and click on it. Consumers are heterogeneous with respect to their receptivity, or preference, for message design, i.e., they differ in terms of what creative elements (tone of the text, color schemes, types of images and their placement on the webpage, etc.) attract their attention and pique their interest. The more closely an ad matches a consumer’s preference in the message space, the more likely s/he will click on it to learn about the product and determine whether they would like to buy it. While a firm can always send a generic ad reflective of its baseline content expertise or its marketers’ initial creative effort, it may wish to personalize ads to improve response rates. This requires correctly identifying consumers’ message-preference locations and subsequently generating ads that take these locations into account.

To accomplish these tasks at scale, firms can leverage AI technologies at their disposal.¹ Specifically, they can employ AI algorithms and models, trained on users’ online behavior and characteristics, as well as any previous engagement with content, to form a prediction regarding which ad creative design will best reflect each consumer’s message preferences. We will refer to this type of AI as “IdAI”, given its role in identifying consumer preferences.² The firm recognizes, however, that due to data and technology limitations, this identification prediction may be inaccurate. Notwithstanding, conditional on the predicted message preferences, the firm employs GenAI to craft a personalized ad to each consumer it wishes to target.³ Once again, due to technology shortcomings, the ad produced may contain distortions, which tend to be more pronounced the farther the intended customized ad is from the firm’s baseline, originally crafted message. Upon receiving an ad, a consumer decides whether to click on it, which is a function of how closely it matches their true message preference location and its expected generative AI distortion. A consumer who clicks then evaluates the product and determines whether to purchase it at the price set by the firm.

Our analysis shows that a monopolist’s optimal strategy results in the probability of clicking on an ad exhibiting an inverted-U shape as a function of the consumer’s message-preference distance

¹While firms could attempt to personalize ads manually, the ability to do so effectively at scale would generally be cost prohibitive without employing AI capabilities.

²This form of AI has also been described as “analytical AI” (see [Grewal et al., 2025](#)).

³As a concrete example, Meta (the parent company of Facebook and Instagram) offers advertisers both types of AI tools. Specifically, its *Advantage+ Detailed Targeting* product uses AI to help marketers identify prospects and understand each individual’s ad responsiveness propensities, while its *Advantage+ Creative Generator* product takes an advertiser-provided creative and uses generative AI to customize the supplied creative (by adjusting image, text, color, and other elements) for each user that will be served with an ad. See [Dolan et al. \(2025\)](#) for more details.

from the firm’s baseline ad generation location. This occurs because the firm is not entirely sure about consumer message preferences, and its GenAI exhibits increasingly greater distortion the farther an intended ad creative is from the baseline message crafted. Hence, the firm chooses to serve consumers who it deems to be relatively distant with ads that are more biased towards the firm’s baseline creative location, resulting in a greater likelihood of a mismatch and a lower click rate. Interestingly, we further find that the firm may target fewer consumers, i.e., advertise less, as its IdAI capability improves, provided the GenAI capability is not too accurate. This occurs because when a firm is initially better able to identify message preferences, and since advertising is costly, it may want to avoid trying to (imperfectly) personalize ads to consumers who are located relatively far away from its baseline creative - as there is still a significant chance of “misfiring” with content that is distorted. However, as IdAI improves further, the misfiring chance declines and the firm is more willing to target extreme-preference consumers, even if its GenAI capability is not perfect.

We then examine duopolistic competition, under the assumption that a consumer is more likely to attend to the ad that appeals to them most, i.e., is closest to their message preference location. We show that, surprisingly, in equilibrium firms’ profits increase with better IdAI capability but decrease with better GenAI capability. The reason for this outcome is that, as the latter capability improves, there is less concern over generative content distortion and firms can better personalize ads to match consumers’ tastes. Hence, each firm is inclined to advertise to consumers who are farther away from its baseline message. But because this behavior results in greater overlap in the firms’ ad targeting strategies, competition is intensified. By contrast, when IdAI capability improves, each firm is inclined to focus on consumers closer to its baseline message location, resulting in less overlap in the consumers targeted and hence reduced competition.

We further explore the implications of letting the firms choose between two discrete levels of GenAI capability that are associated with different costs, for example, different tiers of an AI service. Our analysis reveals that, since investments in GenAI capabilities form strategic substitutes, when the cost of selecting the more advanced GenAI engine is not too high or too low, an asymmetric equilibrium emerges.

Finally, we solve for an extension in which consumers are also heterogeneous in their product preferences, and the firm can imperfectly identify these preferences. Our main findings here are as follows. Price tends to rise, as the firm generally finds it optimal to concentrate on consumers with a higher willingness to pay. However, for a given level of product preference identification accuracy, improvements in the GenAI capability may lead the firm to reduce price. In this case, the improved

ability to personalize ad content makes it optimal to expand advertising outreach to more distant consumers, which requires lowering prices to incentivize purchase.

Taken together, the findings highlight how a firm’s advertising strategy and profits are impacted by two key ingredients of ad personalization at scale – namely, identifying which ad design each consumer would find most resonant to prompt a click and crafting content to achieve a desired design – capabilities that have been transformed by recent AI advances. Moreover, the results reveal that “more” of these AI-enabled capabilities may not always be better. In particular, competing firms may earn greater profits when its GenAI capability is less advanced.

The remainder of the paper proceeds as follows. In the next section, we relate our work to the extant literature on ad targeting and personalization. Subsequently, we formally describe our model structure. This is followed by an analysis of a monopolist that, given the AI capabilities it possesses, chooses its optimal advertising strategy, consisting of which customers to target and what ads to send. We then analyze a duopoly setting, also endogenize the decision of how much to invest in GenAI capability. An extension introduces consumer heterogeneity in the product space, with the firm having imperfect ability to identify consumers’ message and product preferences while using GenAI to personalize ads. We conclude by summarizing our main findings, offering managerial implications, discussing the limitations of our research, and offering opportunities for future research.

2. Literature Review

There is a vast literature in marketing and economics on firms’ incentives to advertise, depending on the impact ads can have on customers’ buying behavior and on the advertising technology or media options available (see, for example, [Bagwell, 2007](#), and references therein). Most relevant to our research is work that has examined the role and implications of ad targeting.

In their seminal paper on the topic, [Iyer et al. \(2005\)](#) analyze a duopoly where firms can target their costly informative ads differentially – to their respective loyal segments as well as to a comparison shopper segment that is willing to buy from either firm depending on price. They find that such targeting can result in greater equilibrium profits, as it allows each firm to focus on customers with a stronger preference for their offering and avoid wasted advertising (to the rival’s loyalists). They further show that endogenizing the ability to target ads to different segments can lead to an asymmetric equilibrium where only one firm invests in this capability.

[Zhang and Katona \(2012\)](#) also examine a three-segment demand structure (with loyalists and “price

shoppers”) where ads can be targeted based on the context in which they appear, thereby exploiting the fact that consumers’ website preferences are often correlated with their product preferences. Firms bid for ad placements on websites through an intermediary that can strategically adjust the accuracy of the targeting technology. They find that when competition among advertisers is intense, in equilibrium, the intermediary is inclined to lower the accuracy.⁴ In a somewhat related paper, [Bergemann and Bonatti \(2011\)](#) examine a scenario whereby different media tend to attract consumers with different product preferences, providing an incentive for firms to advertise in outlets that attract relevant potential buyers. They find that as each consumer segment is more concentrated on specific media (in other words, the targeting accuracy increases) the price of advertising can decline. This is because firms are motivated to reduce overlap in ad placements, thus mitigating competition for ad slots.

Our work differs from these papers in several important ways. First, we distinguish between two aspects of targeting - preference for message (or ad) design and separately for the product. The former, which is absent in the aforementioned papers, is relevant for triggering a response from the consumer (e.g., a click) that subsequently leads to purchase consideration.⁵ Second, and critically, we make a distinction between the ability to accurately identify preferences for message design and the ability to accurately generate ad creatives that match these preferences at the individual level (i.e., ad personalization). Analyzing a firm’s optimal advertising strategy as a function of these two capabilities constitutes a central feature of our study. Notwithstanding, our competitive setting also yields an asymmetric equilibrium in the decision to invest in GenAI capabilities, akin to the finding by [Iyer et al. \(2005\)](#) when the cost of such investment is in an intermediate range. Lastly, we do not include a strategic third-party intermediary (as in [Zhang and Katona, 2012](#)) that serves the ads and controls the accuracy of targeting.

It is also worth mentioning the work by [Grossman and Shapiro \(1984\)](#), who studied the role of advertising in increasing the probability that a consumer finds a product that best matches their taste. They find that as advertising technology becomes more efficient, which can be interpreted as an increase in the ability to target ads to the relevant population interested in the category (see

⁴[Chen et al. \(2001\)](#) as well study a three-segment market, with a focus on individual pricing offers. They allow for imperfect targeting ability, in the sense that customers may be misclassified, and show that the profits of competing firms will initially be higher (win-win) and subsequently lower as the identification accuracy improves. Our work differs in that we concentrate on identification accuracy for message design preferences with consumer sensitivity to a mismatching ad, rather than on price sensitivity.

⁵We point out that in our conceptualization and modeling, ads serve as the initial “gateway” to subsequently learning more about product fit post-click. In that sense, ads in our setting are uninformative about the exact utility the consumer will derive from the product. Our focus is thus on what many marketing scholars and practitioners refer to as the “stopping power” of ads ([Pieters et al., 2010](#)).

discussion in [Grossman and Shapiro, 1984](#), pg. 71), firms’ profits may decrease under some conditions. Our findings are much more nuanced, in that we show that improvements in identification accuracy boost competing firms’ profits while the opposite is true for content personalization accuracy – a hallmark of the recent use of GenAI tools in marketing.

Another research stream in the realm of targeted advertising focuses on what consumers might infer from firms’ ads, conditional on how accurately preferences can be identified. For example, [Shin and Yu \(2021\)](#) study a competitive scenario in which consumers infer their match to the category when receiving an ad, as well as update their beliefs about the quality of the advertiser’s product. In their model, advertising more intensely entails a tradeoff: benefiting from being more prominent vs. prompting consumers to conduct more search, thus allowing the rival to free ride. They find that a higher quality firm has a greater incentive to advertise aggressively and that more precise targeting may result in less advertising.⁶ Our focus is different from [Shin and Yu \(2021\)](#) in that we presume that all consumers are generally interested in the category and have heterogeneous (along a circle) preferences for the product(s). Moreover, we emphasize the benefits to crafting ads that map onto consumers’ heterogeneous tastes for creative appeal, in order to engender an initial interest in the advertiser’s offering and to produce a click. As in [Shin and Yu \(2021\)](#), we also find that as the accuracy of identifying consumer preferences increases, the total amount of advertising can decrease. However, our result is not driven by fears of rival free-riding, but rather due to a desire to avoid sending costly ads that may not hit the mark due to GenAI’s limitations. Additionally, under competition, we show that improvements in different ad targeting aspects - namely identifying consumer message preferences and customizing ads - can have opposing effects on firms’ equilibrium advertising reach and profits.

Lastly, we point out that several works in marketing have highlighted the importance of considering how content, such as ads or web pages, is organized and presented as this can have a measurable impact on consumer responsiveness. For example, [Pieters et al. \(2010\)](#) examine the “stopping power” of ads, i.e., their ability to capture the consumer’s attention in likable ways. They show (using eye tracking) that the creative design (e.g., how irregular and dissimilar the objects in the ad are or how they are arranged in the copy) affects consumer attention. [Hauser et al. \(2009\)](#) demonstrate that individuals prefer website designs that match their ‘cognitive style’ (e.g., analytic, impulsive,

⁶Also relevant noting is work by [Gardete and Bart \(2018\)](#) who look at how a principal, such as a salesperson or advertiser, tailors their communication based on information learned about the potential customer in order to persuade them to take action (such as buy a product). In a cheap talk setting, they show that greater accuracy in identifying preferences can hurt all parties involved because the prospective customer cannot tell if the advertiser is truthful in their statements (or just trying to pander).

or verbal), which in turn can affect conversion rates. While their focus is on how a website might morph and adapt to meet the user’s preferences, they emphasize two challenges that are related to our research: 1) the need to identify a consumer’s preference for website design and 2) the ability to execute on this and produce designs that match the identified cognitive style. In many ways, these correspond to our modeling of AI for identifying consumers’ message design preferences and GenAI’s ability to produce personalized designs at scale given this identification.

To summarize, the contribution of our research is three-fold. First, we introduce in an analytic modeling framework the notion that an ad’s content design can affect the likelihood that an individual will engage with it; and that consumers can be heterogeneous with respect to which design will be most resonant. Second, we break down the ability to effectively target consumers with ads into identifying each consumer’s preference for message design and the ability to then personalize the creative and take this preference into account. Notably, our analysis reveals that improvements in these two capabilities can have differential effects on which consumers are targeted with ads and on firm profits. Third, given our explicit modeling of automated ad customization at scale, to the best of our knowledge, we are among the first to analytically examine the role GenAI can play in firms’ advertising strategies.

3. Monopolist

In this section, we describe our main model setup and solve it for the case of a monopolist. Subsequently, in Section 4, we build upon this structure and introduce competition.

3.1. Model-setup

A monopoly firm operates in a market with a unit mass of consumers. Consumers are heterogeneous with respect to the type of message that catches their attention and prompts them to click on an ad. Specifically, consumers are uniformly distributed on a circle of circumference 2, and consumer i ’s message preference is captured by their location m_i^c .

The firm has expertise in creating a particular style, set of elements, and tone of a message captured by location m^F ; this can also be interpreted as the initial creative the firm’s marketing team had devised. The firm can use Generative AI (GenAI) to modify the initial creative and produce additional messages at scale. That is, the firm can target each consumer with a different message m . We capture how effectively GenAI can generate creative content for messages that differ

from the original message m_F it was trained on with the parameter $\beta \in [0, 1]$. Let $d(m', m'')$ be the shortest distance along the circle between any two locations m' and m'' . Then, if consumer i , located at m_i^c , receives a message m generated by the firm's GenAI, their probability of clicking on the message is

$$Pr[\text{click} \mid m_i^c, m] = 1 - \frac{1}{2}(1 - \beta)d(m^F, m)^2 - \frac{1}{2}d(m_i^c, m)^2. \quad (1)$$

The last term in the above expression captures the attenuation in click probability when an ad differs from the consumer's message preference location; and the greater the difference, measured as the distance $d(m_i^c, m)$ between the ad's creative and the consumer's location, the larger this attenuation. The middle term reflects the attenuation due to the potential distortion of using GenAI to craft individual messages that differ from the baseline creative. Similarly, the greater the distance $d(m^F, m)$ between the message sent and the baseline location the larger this attenuation.⁷ Notably, as $\beta \rightarrow 1$, the GenAI engine is highly capable and this distortion is minimal.

The firm does not ex-ante know consumers' preferences $\{m_i^c\}_{i \in [0, 2]}$ and uses an AI algorithm to identify consumers' message preferences. We denote this algorithm and the underlying data that it uses by IdAI. Specifically, the firm has access to an IdAI-generated database $\{m_i\}_{i \in [0, 2]}$ identifying the message preferences of potential customers. The IdAI database is not perfect and has a precision level α_m : For a mass α_m of consumers drawn uniformly at random from the set of all consumers, IdAI is correct and $m_i = m_i^c$. For the remaining mass $1 - \alpha_m$ of consumers, IdAI is uninformative so that conditional on m_i^c , m_i is uniformly distributed on the circle.

The firm selects which consumers to advertise to and what message to send to each of them. If a consumer clicks on the ad, they are taken to the firm's website, where they learn about the product. They then decide whether to purchase the product. Conditional on clicking on the ad, the expected profit for the firm per consumer is π . Specifically, in the base model, the firm has no information on consumers' product preferences. Therefore, it treats consumers as ex ante homogeneous in terms of their product preferences, and expects the same profit from all consumers who click on an ad. We analyze an extension in which the firm has (imperfect) information on product preferences in Section 5, where the ability to identify consumers' product preferences as well as the firm's pricing play a crucial role. The cost of advertising is c per message sent. We assume that $c < \pi$,

⁷Our use of quadratic functions to represent the attenuation of impact from both sources is similar to the modeling advertising effectiveness in other works. Our results hold qualitatively if instead we were to use an absolute function (though this functional form is less convenient).

ensuring that conditional on clicking on an ad a consumer is expected to be profitable, and note that otherwise the firm never advertises. Table 1 provides a summary of the notation used throughout the paper. All proofs and specific values for cutoff quantities have been relegated to the Appendix.

Table 1: Summary of Model Notation

Notation	Definition
m_i^c	Actual message-preference location of consumer i
m_i	Message-preference location of consumer i as identified by IdAI system
m^F	Baseline message location of the firm (its initial ad creative)
m	Generic message a firm sends
m_i^*	Optimal message a firm sends to a consumer the IdAI indicates is located at m_i
α_m	Accuracy of IdAI technology (probability that $m_i = m_i^c$)
β	Accuracy of GenAI technology
c	Cost of delivering an ad to a consumer
$d(\cdot, \cdot)$	Distance along the circle between any two locations
π	Expected profits from a consumer who clicks on an ad
s	Size of the segment the firm targets with ads
C	Cost of investment in enhanced GenAI capability
x_i^c	Actual product-preference location of consumer i
x_i	Product-preference location of consumer i as identified by product IdAI system
x^F	Product location of the firm
α_p	Accuracy of product IdAI technology (probability that $x_i = x_i^c$)
p	Price of the product

3.2. The Optimal Message for Each Consumer

To find the optimal message the firm would elect to send to any consumer, we first find the probability that consumer i with IdAI location m_i clicks on any message $m \in [0, 2]$. Per equation (1) this probability depends on how close the consumer's true message preference is to the message received, as well as the distance between this message and the firm's baseline message m^F . If IdAI correctly identifies a consumer's message preference (which occurs with probability α_m), the consumer's actual message preference is indeed m_i and its distance $d(m_i^c, m)$ to message m equals $d(m_i, m)$. However, if IdAI incorrectly identifies the consumer's message preference location (which occurs with probability $1 - \alpha_m$), the consumer's message preference is only known to be uniformly distributed over the circle, and its distance $d(m_i^c, m)$ to message m is unknown. Specifically, $d(m_i^c, m)$ conditional on IdAI being incorrect is a random variable denoted by δ , with $\delta \sim U[0, 1]$. Thus, taking into account

the accuracy of IdAI, the probability from (1) of consumer i clicking on message m becomes:

$$\begin{aligned} \Pr[\text{click} \mid m_i, m] &= 1 - \frac{1}{2}(1 - \beta)d(m^F, m)^2 - \frac{\alpha_m}{2}d(m_i, m)^2 - \frac{1 - \alpha_m}{2} \int_0^1 \delta^2 d\delta \\ &= \frac{5 + \alpha_m}{6} - \frac{1}{2}(1 - \beta)d(m^F, m)^2 - \frac{\alpha_m}{2}d(m_i, m)^2. \end{aligned} \quad (2)$$

This probability is maximized for a message m_i^* located between m^F and m_i . Specifically, the distance between the optimal message and the firm's base message is

$$d(m^F, m_i^*) = \frac{\alpha_m}{1 - \beta + \alpha_m} d(m^F, m_i).$$

For example, if we normalize m^F to 0, the optimal message is a weighted average between the consumer's IdAI location m_i and the firm's baseline message $m^F = 0$:


$$m_i^* = \begin{cases} \frac{\alpha_m}{1 - \beta + \alpha_m} m_i, & \text{if } m_i \leq 1 \\ 2 - \frac{\alpha_m}{1 - \beta + \alpha_m} m_i, & \text{if } m_i \geq 1 \end{cases}.$$

Figure 1(a) illustrates that the firm's optimal message is located between where the IdAI places the consumer and its baseline message, depending on both the accuracy of IdAI (α_m) and the effectiveness of GenAI (β). Several observations are worth highlighting from the firm's optimal message strategy. If IdAI returns pure noise (i.e., $\alpha_m = 0$), the firm's ad is received by a random consumer in the message space. In that case, the firm minimizes the distortion due to imperfect GenAI by sending its base message $m_i^* = m^F$. As α_m increases, the probability that the consumer's message preferences is indeed m_i rises, and the firm's optimal message m_i^* gets closer to the IdAI location m_i .

With respect to β , if the firm can generate any modified message without distortion or hallucination that would decrease the probability of a click (i.e., $\beta = 1$), it can perfectly personalize its message to the consumer's IdAI preferences $m_i^* = m_i$. However, if the firm's GenAI technology is imperfect ($\beta < 1$), producing ads distant from its baseline creative entails a decreased click probability. Thus, the firm prefers not to attempt to cater fully to the consumer's presumed preference location, and instead sends a message that is closer to m^F (i.e., $d(m^F, m_i^*) < d(m^F, m_i)$). As β increases, the concern over ad distortion due to imperfect GenAI is mitigated and the optimal message m_i^* shifts closer to the consumer's IdAI location m_i .

Plugging in the firm's solution to the optimal message (m_i^*) into the click probability expression

Figure 1: Optimal Message to a Given Consumer (left) and Optimal Segment to Target with Ads (right)



(a) Optimal message m_i^* for a consumer whose IdAI location is identified as m_i

(b) The firm targets a segment of consumers symmetrically around m^F

Note: m^F is normalized to 0 in the expression for m_i^* and $m_i \in [0, 2)$. The subfigure (a) on the left displays the monopolist's optimal message m_i^* for IdAI preference m_i (with probability α_m this is in fact the consumer's preference location m_i^c); otherwise the output from IdAI is uninformative and the consumer can be located anywhere along the circle with uniform probability). The subfigure (b) on the right illustrates the targeted segment centered around m^F .

in (1), yields that the probability of consumer i , with IdAI location m_i , clicking on ad m_i^* is:

$$\Pr[\text{click} \mid m_i, m_i^*] = \frac{5 + \alpha_m}{6} - \frac{\alpha_m(1 - \beta)}{2(1 - \beta + \alpha_m)} d(m^F, m_i)^2. \quad (3)$$

The following lemma characterizes the effect of increasing the IdAI and GenAI capabilities on this probability.

Lemma 1 *For a consumer who is identified to be located at m_i and to whom the firm sends an optimal ad m_i^* , the probability of a click will*

- (i) *initially decrease and subsequently increase as the IdAI capability (α_m) increases, if m_i is sufficiently distant from the firm's baseline ad creative m^F ; and otherwise, the click probability will increase in α_m .*
- (ii) *always increase as the GenAI capability (β) increases.*

One would have expected that as the ability to identify consumers improves, i.e., α_m increases, the firm would personalize its messages such that all consumers will have higher click probabilities. However, as Lemma 1 reveals, this is not the case for consumers presumed to be located not sufficiently close to the firm's baseline ad creative. To gain intuition for this result, it is helpful to refer to expression (2) and consider the impact of increasing α_m on the probability of a click,

while keeping the message itself fixed.⁸ Effectively, an increase in α_m implies that the probability of ad m_i^* reaching a consumer with message preference m_i increases vis-à-vis the probability of this ad reaching a consumer with a randomly drawn message preference. When α_m is small and m_i is sufficiently distant from the firm's baseline message m^F , the targeted ad m_i^* is quite different from m_i . Hence, the fact that this ad is more likely to reach consumer m_i as IdAI initially improves is detrimental to the probability of a click. On the other hand, when α_m is much larger or m_i is very close to the firm's baseline message m^F , then the optimal message the firm personalizes, m_i^* , is already relatively close to m_i . Consequently, as the IdAI capability improves further, the chances that ad m_i^* in fact reaches a consumer with message preference m_i rises and this has a positive impact on the click probability.

Turning to Part (ii) of the lemma, an increase in β improves the accuracy of message customization. When IdAI is not entirely uninformative ($\alpha_m > 0$), more effective GenAI capability allows crafting messages that are closer to consumers' ideal points. This unequivocally increases the probability of a click.⁹

3.3. Targeting Decision

Having established the optimal message a firm would send to each consumer based on identified location m_i and the resulting click through rate it engenders, we now examine how the firm chooses which consumers to target with its ads. When the firm sends ad m_i^* to consumer i with IdAI message preference m_i , the consumer clicks on it with probability $\Pr[\text{click} \mid m_i, m_i^*]$ as given in expression (3). Then, conditional on clicking, the expected profit is π . That is, the expected profit from targeting consumer i is $\Pr[\text{click} \mid m_i, m_i^*]\pi$, and the firm targets m_i if and only if the expected revenue is larger than the cost c of sending an ad, or $\left(\frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)}d(m^F, m_i)^2\right)\pi \geq c$. We note that the LHS of this expression decreases in $d(m^F, m_i)$, the distance between the baseline message the firm produces and the identified location of consumer i . As a result, the firm will target a segment of consumers that is symmetric around the baseline message location m^F .

Let $s \in [0, 1]$ denote the size of the target segment, corresponding to all the consumers with message preferences within a distance s of m^F on the circle (i.e., target all m_i such that $d(m^F, m_i) \leq s$). Figure 1(b) illustrates which consumers are part of the targeted segment, and which consumers

⁸By the envelope theorem, this exercise captures the first order effects of varying α_m . In other words, the change in the optimal message due to a change in α_m is second order compared to instances where the variable α_m itself appears in the probability expression.

⁹Trivially, if $\alpha_m = 0$ or $m_i = m^F$, changes in β do not affect the click probability per (3).

are not targeted. We first find the expected number of clicks $\Phi(s)$ by integrating $\Pr[\text{click} \mid m_i, m_i^*]$, the probability of clicking given in (3), over all the consumers with $d(m^F, m_i) \in [0, s]$, which is

$$\Phi(s) = \frac{5 + \alpha_m}{6}s - \frac{\alpha_m(1 - \beta)}{6(1 - \beta + \alpha_m)}s^3. \quad (4)$$

We then obtain the expected profit from targeting a segment of size s by multiplying the expected number of clicks $\Phi(s)$ by the profit per clicking consumer π and deducting the advertising cost cs , which leads to

$$\Pi = \left(\frac{5 + \alpha_m}{6}s - \frac{\alpha_m(1 - \beta)}{6(1 - \beta + \alpha_m)}s^3 \right) \pi - cs. \quad (5)$$

The firm selects the size of its target segment to maximize its expected profit: $s^* = \arg\max_s \Pi$. The following proposition characterizes whether the firm targets the entire market, part of the market, or does not advertise altogether.

Proposition 1 *There exist thresholds for IdAI and GenAI capabilities $\bar{\alpha}_m, \bar{\beta}$, such that the firm's optimal ad targeting strategy s^* is as follows,*

- (i) *When IdAI accuracy is high ($\alpha_m > \bar{\alpha}_m$),*
 - (i) *if GenAI capability is low ($\beta < \bar{\beta}$), the firm targets a subset of the market ($0 < s^* < 1$),*
 - (ii) *if GenAI capability is high ($\beta > \bar{\beta}$), it targets the entire market ($s^* = 1$);*
- (ii) *When IdAI accuracy is low ($\alpha_m < \bar{\alpha}_m$), if the advertising cost–profit ratio (c/π) is high, the firm does not advertise to any consumer ($s^* = 0$), while if this ratio is low, it targets the entire market ($s^* = 1$).*


When IdAI accuracy is high, the firm finds it optimal to advertise. If its GenAI capability is limited, the firm anticipates substantial message distortion for consumers located far from the baseline message, which lowers their click probability. As a result, the firm targets only a subset of consumers, and the optimal solution involves less than full market coverage. By contrast, when GenAI capability is sufficiently strong, click probabilities are high enough to induce full market coverage.

When IdAI accuracy is low, however, the firm's advertising decision depends solely on the advertising cost–profit ratio c/π . In this case, the information about consumers' message preferences is too noisy to support meaningful targeting. In particular, in the extreme case of $\alpha_m = 0$, all consumers appear identical from the firm's perspective. It is therefore optimal either to advertise to all consumers or to none. If α_m is sufficiently low and c/π is high, the firm refrains from advertising, since even consumers believed to be close to the baseline message m^F exhibit too low an expected

purchase probability to justify the advertising cost. Conversely, if c/π is low, the firm advertises to the entire market.

Figure 2 provides a visual illustration of Proposition 1 for two different values of c/π , the advertising cost relative to expected profit per clicking consumer.

Figure 2: Target Size (s^*) as a Function of IdAI (α_m) and GenAI (β) Capabilities



Note: When $\frac{c}{\pi}$ is very low, $s^* = 1$ for any IdAI (α_m) and GenAI (β) capabilities.

The effect of c/π on targeting is intuitive. As the cost of sending an ad increases relative to the potential profit, the firm requires a higher expected click probability in order to target a consumer. From (3), expected click probability decreases with distance from the baseline message m^F . Consequently, as c/π rises, the size of the targeted segment shrinks.

The next proposition focuses on the region where the firm targets a subset of consumers, i.e., $s^* \in (0, 1)$. It completes the characterization of the firm's targeting decision by deriving the impact of changes to the GenAI and of IdAI capabilities on the size of the targeted segment and the firm's profit.

Proposition 2 *If GenAI capability is such that the firm optimally targets only a subset of consumers with its ad $s^* \in (0, 1)$, the targeted segment size*

(i) *increases in the GenAI capability β ;*

(ii) *initially decreases and subsequently increases as the IdAI capability (α_m) increases iff the advertising cost-profit ratio is not too high; otherwise, the segment size will increase in α_m .*


Moreover, the profit of the monopolist weakly increases in both the effectiveness of GenAI and the

accuracy of IdAI.

The intuition for the first part of Proposition 2 mirrors Part (ii) of Lemma 1: an improvement in the GenAI capability increases the probability that any consumer would click on the optimal ad the firm would personalize for them. This makes it unequivocally more worthwhile to target additional consumers as GenAI improves.

The second part of Proposition 2 once again reveals the possibility of an intriguing non-monotonic effect for IdAI. In particular, there exist conditions for the firm to contract the size of the targeted segment as this capability improves. Figure 3 illustrates how the size of the targeted segment depends on the value of IdAI for two advertising cost–profit ratios. The intuition for the non-monotonic effect of IdAI technology improvement is as follows. When the advertising cost–profit ratio is not too high, the firm is inclined to advertise to more consumers when it can better identify message preferences. However, recall from Part (i) of Lemma 1 that initially as α_m increases the click through rate of some consumers, specifically those distant from m_F , declines such that $\Pr[\text{click} \mid m_i, m_i^*] < c/\pi$. Hence, the firm optimally forgoes additional distant consumers as IdAI improves because they become unprofitable to target in expectation. Said differently, as IdAI initially improves the firm prefers to focus on consumers whose message preferences are closer to the baseline ad creative, since their true preferences (m_i^c) is more likely to correspond to what the AI has identified m_i and the GenAI distortion is minimized.

Figure 3: Impact of the Accuracy of IdAI (α_m) on the Size of the Target Segment (s^*)



4. Competition

We now examine duopolistic competition within the context of our model. Consider two firms, denoted as Firm 0 and Firm 1. Let Firm j 's baseline message creative be located at m_j^F . The firms are situated on opposite sides of the circle from each other, and, without loss of generality, we set $m_0^F = 0$ and $m_1^F = 1$.¹⁰ For simplicity, and to focus on the competitive forces relevant in deciding how to personalize ads and whom to target, we assume that if presented with at least one ad, the consumer always clicks on exactly one ad. This implies that if a consumer receives ads from both firms, the consumer will attend to only one of them.¹¹

We now describe how a consumer presented with ads from both firms determines which ad to click on. Consistent with the basic model setup, we let the attractiveness of an ad by firm j with message creative m for consumer i be: $V_j(m, m_i^c) = 1 - \frac{1}{2}(1 - \beta_j)d(m, m_j^F)^2 - \frac{1}{2}d(m, m_i^c)^2$. A consumer, in general, attends to the ad that is more attractive to them. However, on any given occasion, they may experience idiosyncrasies that affect their attention span, captured by a stochastic component ϵ_i distributed $U[-1, 1]$ i.i.d. across consumers and locations.¹² Taken together, the consumer clicks Firm 0's ad with probability

$$Pr[\text{click Firm 0} \mid m_i^c, m_0^F, m_1^F] = Pr[V_0(\cdot) - V_1(\cdot) > \epsilon_i] \quad (6)$$

$$= \frac{1 + V_0(\cdot) - V_1(\cdot)}{2}, \quad (7)$$

and clicks Firm 1's ad with the remaining probability. We denote by $m_{j,i}^*(\cdot)$ the optimal message that Firm j sends to consumer i in equilibrium.

We maintain the same roles of GenAI and IdAI technology from the monopoly setup and further assume that the firms use the same database to calibrate their IdAI algorithm. To be specific, for every consumer i both firms have access to the same IdAI consumer message preference m_i . We initially assume that firms possess similar GenAI capabilities, though we later allow them to endogenously access different GenAI levels, denoted β_j for firm $j \in \{0, 1\}$.

¹⁰All the results presented would hold qualitatively if the firms are placed at any non-identical locations along the circle. We only require that there be some overlap in their ad targeted segments in equilibrium.

¹¹Our results would qualitatively hold if the consumer also clicks on a second, less attractive, with a lower probability than the more attractive ad. This analysis is available from the authors upon request.

¹²We note that this formulation defines a continuum of random variables and requires defining appropriate probability spaces for consumers, as well as modeling uncertainty to invoke the Exact Law of Large Numbers and address measurability issues. The construction is available in Sun (2006) and Sun and Zhang (2009), and summarized in Appendix C1 of Fainmesser et al. (2023).

In analogy to the monopoly case, conditional on a consumer clicking on firm j 's ad, the expected profit for firm j from that consumer is π .¹³

4.1. Timeline and equilibrium

The timeline of the competition model is as follows:

- (i) Firms simultaneously choose which consumers to target and which message to send to each targeted consumer (at a cost c per ad sent).
- (ii) The ads are displayed and each targeted consumer clicks on one ad as described above.
- (iii) Sales and profits are realized.

In Section 4.4, we enrich this model by adding a stage 0 in which firms decide whether to invest in GenAI capability enhancements and study their investment decisions. Our solution concept throughout is the Subgame Perfect Nash Equilibrium, henceforth equilibrium.

Notably, each firm's decision problem can be broken into two parts that can be solved separately. Part 1: the firms decide what message to send to each targeted consumer who will receive both ads, and Part 2: the firms decide which consumers to target. In a sense, Part 1 captures the competition over each consumer, whereas Part 2 captures the firms' decisions of which consumers to compete on.

4.2. Messaging: competing for consumer i 's Click

In making the decision regarding which message to send to a given consumer i , each firm maximizes the attractiveness of its ad to that consumer. To achieve that, Firm 0's optimal message, $m_{0,i}^*$, is the same as m_i^* in the monopoly case, and Firm 1's message can be derived symmetrically based on Firm 1's baseline message of $m_1^F = 1$.

Noting that, we begin by analyzing a benchmark case in which both firms are using the same exogenously given GenAI technology, i.e. $\beta_0 = \beta_1 = \beta$ for some $\beta \in [0, 1]$.

Lemma 2 *In equilibrium, the probability that consumer i , who is presented with ads from both firms, clicks on firm j 's ad*

- (i) *decreases in the consumer's IdAI message preference distance from firm j (i.e., $d(m_j^F, m_i)$)*
- (ii) *increases in IdAI accuracy α_m if and only if the consumer's identified message preference is closer to firm j than to its competitor.*
- (iii) *increases in the GenAI capability β if and only if the consumer's identified location is farther*

¹³This captures the idea that the firms do not know ex-ante consumers' valuation for each firm's product. In Section 5, we explicitly model consumers' product preferences.

from firm j than from its competitor.

Lemma 2, reveals a number of noteworthy properties of the click rate in equilibrium. Part (i) of the lemma is intuitive and corresponds to the same result as in the monopoly case. Specifically, the farther a consumer is from a firm's baseline ad creative position, the less likely they are to receive an ad that matches their ideal location - due to both imperfect identification and imperfect GenAI message personalization.

Parts (ii) and (iii) of Lemma 2 capture more of the richness of the competitive interaction between the firms in the presence of IdAI and GenAI. On the one hand, as IdAI capabilities improve (α_m increases), it becomes more advantageous to craft a message that corresponds to the consumer's IdAI message preference location. This is because the consumer's location, as identified by the IdAI system, is more likely to match the consumer's *actual* message preference location (i.e., m_i is more likely to equal m_i^c). On the other hand, when GenAI capabilities improve (β increases), a firm's relative advantage at crafting ads to consumers presumed close to its baseline location diminishes. This is because the improvement in the rival firm's ability to tailor ads to far away consumers as β increases is more pronounced.¹⁴

4.3. Targeting Decisions

If the cost of advertising, c , is sufficiently large, then for any consumer i , regardless of her message preference, at most one firm will find it profitable to display an ad to that consumer in equilibrium. Instead, to allow for meaningful competitive interaction, we focus on a sufficiently small advertising cost such that there exist some (non-zero measure of) consumers to whom both firms would like to advertise. In this case, the equilibrium has a unique structure: there exist s_0, s_1 such that $s_1^* + s_0^* \geq 1$, Firm 0 displays an ad to consumer i if and only if $d(m_i, 0) \leq s_0^*$, and Firm 1 displays an ad to consumer i if and only iff $d(m_i, 1) \leq s_1^*$. This structure is illustrated in Figure 4.

The following lemma characterizes the equilibrium structure and provides key comparative statics with respect to changes in the AI capabilities.


Lemma 3 *The equilibrium is symmetric, with both firms choosing the same advertising segment size s^* . That is, $s_0^* = s_1^* = s^*$. The equilibrium advertising segment size s^**

(i) *decreases in IdAI capabilities;*

(ii) *increases in GenAI capabilities.*

¹⁴In the limit, when GenAI capabilities are perfect ($\beta = 1$), a firm's ability to target a consumer is independent of the firm's location and the probability that a consumer clicks on an ad from either firm converges to $1/2$ regardless of the consumer's message preference.

Figure 4: Duopoly Ad Targeting Equilibrium Structure



Somewhat surprisingly, Lemma 3 shows that IdAI and GenAI capabilities have opposing effects on the equilibrium size of the advertising segments of the firms. The intuition for this result follows from Lemma 2. Specifically, an improvement in IdAI tells firm j that the consumer who is identified as being farthest from it among those it advertises to is more likely to actually be farthest from it (i.e., m_i is more likely to correspond to m_i^c). As a result, consumer i is less likely to click on firm j 's ad, and advertising to her may not be a good investment for firm j . In contrast, an improvement in GenAI implies that a firm's personalized ad creative is less likely to “miss the mark”. Hence, the firm suffers a smaller “penalty” when advertising to a consumer with a message preference that is farther from the firm's baseline ad location. Advertising to a previously marginal consumer may now be a strictly profitable investment, and the new marginal consumer is then farther from the firm. This, in turn, increases the set of consumers who receive ads from both firms and intensifies competition for their clicks.

The forces captured in Lemmas 2 and 3 have significant implications for the impact of IdAI and GenAI on firms' profits, as shown in the proposition below.

Proposition 3 *In equilibrium, both firms' profits*

- (i) *increase in IdAI capabilities;*
- (ii) *decrease in GenAI capabilities.*

Proposition 3 reveals a stark contrast with the monopoly case. Recall from Proposition 2 that a monopolist's profits always increase when either the IdAI or GenAI capability improves. However, under competition, while each firm's profits are increasing in the IdAI capability they are decreasing

in the GenAI technology. The reason for this outcome is as follows. An increase in the IdAI capability relaxes competition by driving each firm to focus on consumers closer to its baseline ad creative; thus reducing the advertising segment sizes (per Lemma 3) and hence the set of consumers receiving ads from both firms. In contrast, an increase in GenAI capabilities induces each firm to expand its target audience and step further into its “rival’s turf.” In other words, better GenAI capability intensifies advertising competition by prompting each firm to reach into the segment closer to its competitor, which ultimately reduces profits for both.¹⁵ The findings gleaned from Proposition 3 and in Lemmas 2 and 3 are summarized in Table 2.

Table 2: The Effects of Improvements in IdAI and GenAI (Competition)

	Probability of click near m_j^F	Probability of click far from m_j^F	Advertising interval s_j^*	Profits
Improved IdAI ($\alpha \uparrow$)	$\uparrow$	$\downarrow$	$\downarrow$	$\uparrow$
Improved GenAI ($\beta \uparrow$)	$\downarrow$	$\uparrow$	$\uparrow$	$\downarrow$

4.4. Investment in GenAI

Assume now that before advertising decisions are made, each firm can simultaneously decide whether to invest in enhancing its GenAI capabilities or to (costlessly) use basic available technology. In particular, if Firm j does not invest in enhancing GenAI capabilities then $\beta_j = \underline{\beta} < 1$ and if it invests, it pays a fixed cost C and its capability is elevated to $\beta_j = 1$.¹⁶

Solving for equilibrium investment decisions requires extending our analysis above to allow firms to have different GenAI technologies, i.e., derive targeting decisions and profits for the case that $\beta_0 \neq \beta_1$. Details of this analysis are provided in Appendix B.

We are now ready to state our result on firms’ equilibrium investment in GenAI in a competitive setting.

Proposition 4 *The firms’ GenAI investment decisions, β_0 and β_1 , are strategic substitutes. Consequently, there exist thresholds $\underline{C} \leq \bar{C}$ such that in equilibrium*

- (i) *if $C < \underline{C}$, both firms invest in enhanced GenAI capability;*
- (ii) *if $\underline{C} < C < \bar{C}$, exactly one firm invests in enhanced GenAI capability;*
- (iii) *if $C > \bar{C}$, none of the firms invests in enhanced GenAI capability.*

¹⁵Recall that when a firm advertises to the same consumer as its rival does, it incurs the full cost of delivering the ad c but induces a click with a probability that is less than one (given the rival may obtain the consumer’s click with a positive probability).

¹⁶All of our results immediately extend to any sufficiently enhanced capabilities upon investment.

Proposition 4 shows that two, otherwise ex-ante symmetric, firms may end up in an asymmetric equilibrium whereby one firm enhances its GenAI capabilities while its rival sticks to a more basic GenAI technology. For the proposition to hold, it must be the case that the increase in a firm's profit following investment in GenAI is greater if the other firm did not invest in GenAI (than if the other firm did invest in GenAI). We indeed establish that this is the case in the proof of Proposition 4 (see Appendix).

To understand the intuition, consider the following. If a firm enhances its GenAI capability, and thus achieves $\beta = 1$, it is able to accurately personalize messages without any distortion. Consequently, and irrespective of its rival's actions, the firm will advertise to all consumers.¹⁷ Next, start from the case where both firms do not invest in enhancing their GenAI and suppose that their (symmetric) targeted advertising interval is some $s_{\underline{\beta}}^* < 1$. If Firm 0 then deviates and invests in GenAI, given the observation above, it will seek to cover the entire market, i.e., it increases its target segment size by $1 - s_{\underline{\beta}}^*$ at a cost of $c(1 - s_{\underline{\beta}}^*)$. Given this move by Firm 0, Firm 1 will find it optimal to reduce its advertising interval, say to some $s_{\underline{\beta},1}^* < s_{\underline{\beta}}^*$. If, however, Firm 1 also decides to invest in enhancing its GenAI capability, it would as well wish to cover the entire market; and thus increase its target advertising segment size by $1 - s_{\underline{\beta},1}^*$ at a cost of $c(1 - s_{\underline{\beta},1}^*) > c(1 - s_{\underline{\beta}}^*)$. Then, it is left to note that this higher advertising cost does not give Firm 1 any greater sales boost than the boost Firm 0 initially received. Thus, the overall net profit gain for Firm 1 from following suit is lower than that of Firm 0, and we can always find a cost of GenAI enhancement (C) for which Firm 1 will not find it profitable to incur this investment.

5. Heterogeneous product preferences

In this section, we analyze a monopolist that may have access to some information about consumers' product preferences in addition to message preferences. Specifically, consumers are heterogeneous with respect to message preferences m_i^c and product preferences x_i^c . As in the main model, message preferences are uniformly distributed on a circle of circumference 2 (message space). Product preferences are independent of message preferences and uniformly distributed on a circle of circumference 2 (product space).

Let the firm sell a product located at x^F on the product-preference circle at price p . A consumer

¹⁷Recall that we focus on the parameter region where firms will always have an incentive to advertise to at least some mutual consumers, i.e., there will be a non-zero overlapping segment targeted by both firms.

located at x_i^c in the product space, who buys the product, obtains utility

$$v - td(x^F, x_i^c) - p,$$

where v denotes baseline utility and t measures the marginal disutility from product–consumer mismatch, captured by the distance $d(x^F, x_i^c)$ in the product space. The outside option is normalized to 0, so a consumer who clicks on an ad will purchase if and only if $d(x^F, x_i^c) \leq (v - p)/t$.

As before, the firm does not observe $\{m_i^c, x_i^c\}_i$ and relies on a database $\{m_i, x_i\}_i$ generated by IdAI algorithms. The identification processes of the two aspects of a consumer’s preferences are independent and imperfect. For message preferences, with probability α_m , identification is correct and $m_i = m_i^c$; with probability $1 - \alpha_m$, the algorithm is uninformative so that, conditional on the true preference m_i^c , the IdAI reported preference m_i is uniformly distributed on the message circle. Similarly, for product preferences, with probability α_p , identification is correct and $x_i = x_i^c$; with probability $1 - \alpha_p$, conditional on x_i^c , x_i is uniformly distributed on the product circle.

The main model in Section 3 corresponds to the case in which $\alpha_p = 0$. That is the firm has no information about product preferences. In that case, conditional on clicking, a consumer purchases with probability $(v - p)/t$ for $p \in [v - t, v]$, implying that the expected profit conditional on clicking is $\pi = \frac{v-p}{t}p$. We assume $v/2 \leq t \leq v$, and so π is maximized at $p = v/2$ over $[v - t, v]$. From the main model, we know that the firm’s expected profit strictly increases in π , hence the firm sets $p^* = v/2$ when $\alpha_p = 0$ and targets only based on the message preferences, with the optimal size of the target segment in the message space s^* given by Proposition 1 applied to $\pi^* = v^2/(4t)$.

When $\alpha_p > 0$, the purchase probability depends on the identified product location x_i . If $d(x^F, x_i) \leq (v - p)/t$, the consumer buys with probability 1 when product-preference identification is correct (probability α_p) and with probability $(v - p)/t$ when identification is uninformative (probability $1 - \alpha_p$), implying $\Pr[\text{buy} \mid \text{click}, d(x^F, x_i) \leq (v - p)/t] = \alpha_p + (1 - \alpha_p)\frac{v-p}{t}$. In contract, if $d(x^F, x_i) > (v - p)/t$, the consumer buys only when product-IdAI is uninformative, so $\Pr[\text{buy} \mid \text{click}, d(x^F, x_i) > (v - p)/t] = (1 - \alpha_p)\frac{v-p}{t}$. Thus, the price p set by the firm endogenously partitions consumers into two segments in the product space. The *near segment*, defined by $d(x^F, x_i) \leq (v - p)/t$, generates expected profit conditional on clicking

$$\pi_n = \left(\alpha_p + (1 - \alpha_p)\frac{v-p}{t} \right) p,$$

while the *far segment*, defined by $d(x^F, x_i) > (v - p)/t$, generates

$$\pi_f = (1 - \alpha_p) \frac{v - p}{t} p.$$

Notice that consumers in the near segment have a higher expected profit conditional on clicking than consumers in the far segment: $\pi_n > \pi_f$. Also notice that the sizes of the near and far segments are endogenously determined by the price p . Specifically, the number of consumers in the near segment is $(v - p)/t$, while the number in the far segment is $1 - (v - p)/t$.

For any given price p , the firm chooses which consumers to target in the product and message spaces, as well as the message sent to each of them. Messages are generated according to the GenAI capability and are identical to those in the main model. Recall that targeting in the product space depends on expected profit conditional on clicking, which is either π_n or π_f . Let s_n and s_f denote the optimal sizes of the target segments in message space corresponding to π_n and π_f , respectively. By Proposition 1, the optimal targeting strategy has the following structure. Consumers in the near segment in the product space are targeted if their message preference lies within distance s_n of m^F in the message space: target consumers (m_i, x_i) such that $d(x^F, x_i) \leq (v - p)/t$ and $d(m^F, m_i) \leq s_n$. Consumers in the far segment are targeted if their message preference lies within distance s_f of m^F : target consumers (m_i, x_i) such that $d(x^F, x_i) > (v - p)/t$ and $d(m^F, m_i) \leq s_f$. Since consumers in the near segment have a higher expected profits conditional on clicking than consumers in the far segment ($\pi_f < \pi_n$ when $\alpha_p > 0$), when the firm advertises, it targets more consumers in the near than in the far segment $s_f < s_n$ or it covers the entire market $s_f = s_n = 1$.

Finally, the firm selects the price p to maximize its overall expected profit, which is¹⁸

$$\left[\Phi(s_n^*) \left(\alpha_p + (1 - \alpha_p) \frac{v - p}{t} \right) p - cs_n^* \right] \frac{v - p}{t} + \left[\Phi(s_f^*) (1 - \alpha_p) \frac{v - p}{t} p - cs_f^* \right] \left(1 - \frac{v - p}{t} \right).$$

The optimal price p^* is associated with the optimal targeting policy, yielding segment-specific message reach s_n^* and s_f^* in the message space for the near and far product segments, respectively. In the rest of the section, we focus on the case when $s_n^* > 0$, i.e., it is optimal for the firm to advertise to some consumers.

The following proposition compares the optimal price p^* to the benchmark price $v/2$ that arises when the firm has no information about product preferences, i.e., when $\alpha_p = 0$.

¹⁸Per (4), $\Phi(s)$ is the expected number of clicks when targeting a segment of size s in the message space.

Proposition 5 *The optimal price satisfies:*

- (i) $p^* > \frac{v}{2}$ iff $\alpha_p > 0$ and $s_f^* < 1$;
- (ii) $p^* = \frac{v}{2}$ otherwise.

This proposition shows that, when the firm has some information about product preferences ($\alpha_p > 0$) and does not advertise to the entire market ($s_f^* < 1$), the optimal price exceeds the price charged in the absence of product-preference information. With product-IdAI, the firm can identify consumers with higher expected willingness to pay and concentrate its advertising effort on them, thereby sustaining a higher price. The mechanism, however, is more nuanced than simply identifying high-willingness-to-pay consumers.

To build intuition, consider the case when the advertising cost c is sufficiently low that the firm advertises to the entire market when $\alpha_p = 0$ and to the entire near segment when $\alpha_p = 1$. When $\alpha_p = 0$, the firm chooses p to maximize $\Phi(1)p(v - p)/t - c$. The optimal price therefore maximizes expected profit conditional on clicking and equals $v/2$, independent of advertising considerations. In contrast, when $\alpha_p = 1$, the firm never advertises to the far segment because those consumers never purchase after clicking. Therefore, the firm chooses p to maximize $[\Phi(1)p - c](v - p)/t$, which internalizes both the advertising cost c and the expected number of clicks $\Phi(1)$. Therefore, $p^* = \frac{v+c/\Phi(1)}{2} > \frac{v}{2}$. For both $\alpha_p = 0$ and $\alpha_p = 1$, a higher price reduces the number of customers who buy conditional on clicking, which is related to $(v - p)/t$. The key difference is that when $\alpha_p = 1$, price affects not only margins and demand conditional on clicking, but also the size of the near segment and therefore the number of consumers targeted. A higher price reduces the number of consumers to whom the firm advertises (since we assume in this example that $s_n^* = 1$). As a result, advertising expenditures become price-dependent rather than fixed. The firm thus trades off a smaller customer base against both higher margins and lower advertising expenditures. Because advertising costs effectively enter the marginal cost of serving additional consumers, the optimal price is shifted upward relative to $v/2$. Put differently, targeting with product information links pricing and advertising decisions, whereas without product information the two are separable.

We next characterize how GenAI capability β affects optimal pricing and targeting decisions.


Proposition 6 *The impact of GenAI capability β on price and targeting decisions is characterized as follows:*

- (i) *If $s_n^* < 1$, then p^* is independent of β ; if $s_n^* = 1$, then p^* strictly decreases in β .*
- (ii) *Both s_n^* and s_f^* increase in β .*

Recall that GenAI capability β increases the effectiveness of messaging by raising the probability that a targeted consumer clicks. This proposition shows that if the firm can still expand advertising reach in the near segment (i.e., $s_n^* < 1$), the optimal price remains unchanged as β increases, whereas if the firm already advertises to the entire near segment ($s_n^* = 1$), the optimal price falls with β .

For simplicity, we explain the intuition when it is optimal to advertise only to the near segment, i.e., $s_f^* = 0$. An increase in β generates two effects. First, a higher β expands s_n^* per Proposition 2, so the firm benefits from raising the expected profit per clicking consumer, π_n , which can be achieved by increasing p^* . Second, a higher β improves the efficiency of generating clicks, increasing the expected number of clicks $\Phi(s_n^*)$ from those consumers who are already targeted. Thus, the firm benefits from raising the total expected profit from the near segment $\pi_n(v - p)/t$, by increasing the size of that segment, which can be accomplished by reducing p^* . When $s_n^* < 1$, these two effects exactly offset each other, leaving p^* unchanged. If s_n^* is already maximized at 1, only the second effect operates, and the firm lowers p^* in response.

Figure 5: Impact of GenAI capability on the optimal price and advertising target decisions (when $0 < s_f^* < 1$)



(a) When $s_n^* < 1$, as GenAI capability β improves, the firm does not change its optimal price p^* and advertises to more customers in both the near and far segments.

(b) When $s_n^* = 1$, as GenAI capability β improves, the firm lowers its optimal price p^* , which augments the size of the near segment in the product space, and also advertises to more customers in the far segment.

Figure 5 illustrates how improved GenAI capability affects the firm’s pricing and advertising strategies when it targets both product-space segments ($0 < s_f^*$). In the left panel, the firm has room to expand message reach in the near segment, so the optimal price remains unchanged as β increases, while message reach expands in both the near and far segments. In the right panel, the firm already advertises to the entire near segment. In this case, an increase in β leads to a lower price (effectively growing the near segment) and is accompanied by increased advertising in the far segment.

Finally, if the GenAI capability of a monopolist improves, the firm derives a higher profit and consumer surplus also increases, since the price is weakly decreasing and more consumers have the opportunity to consider purchasing buying the product.

6. Conclusion

Recent advances in Generative AI (GenAI) technology have fundamentally altered firms’ ability to personalize marketing communications at scale. In this paper, we develop an analytical framework to examine how two distinct AI-enabled capabilities—(i) the ability to identify consumers’ message preferences (IdAI) and (ii) the ability to accurately generate customized ad content (GenAI)—shape firms’ advertising strategies, competitive outcomes, and pricing decisions.

Our analysis yields several key insights. First, in a monopoly setting, improvements in GenAI capability strictly increase advertising reach and profits by reducing distortion in personalized content. By contrast, improvements in message preference identification (IdAI) can have non-monotonic effects on both click-through rates and the size of the targeted segment. When generative capability remains imperfect, better identification may initially cause the firm to contract its advertising reach, as it seeks to avoid costly attempts to personalize ads to consumers whose inferred preferences lie far from its baseline creative. Thus, more accurate identification does not always imply broader advertising.

Second, under competition, the effects of the two AI capabilities diverge sharply. Improvements in IdAI relax competition by inducing firms to concentrate on consumers closer to their respective baseline creative positions, thereby reducing overlap in advertising coverage and increasing equilibrium profits. In contrast, improvements in GenAI capability intensify competition: as message distortion concerns diminish, each firm optimally expands its advertising reach into segments closer to its rival, increasing overlap in consumers targeted and reducing profits. Hence, unlike in monopoly,

the availability of “better” generative capability can actually reduce firm profitability in competitive markets.

Third, when GenAI investment is endogenous, capabilities become strategic substitutes. For intermediate investment costs, equilibrium asymmetry arises: one firm upgrades to advanced GenAI capability while the rival does not. This highlights that AI investment incentives depend critically on rivals’ technological choices.

Finally, extending the model to allow for heterogeneity in product preferences introduces a novel interaction between AI capabilities and pricing. The ability to identify product-preference heterogeneity generally leads to higher prices, as the firm optimally concentrates on consumers with higher willingness to pay. However, conditional on product preference identification accuracy, improvements in GenAI precision may induce price reductions when the firm already covers its core high-match segment. In this regime, improved content personalization makes it optimal to expand outreach to more distant consumers, which requires lowering prices to maintain purchase incentives. Thus, improved advertising AI capabilities can either increase or decrease prices depending on market coverage conditions.

6.1. Managerial Implications

Our findings offer several implications for managers considering utilizing and investing in AI-enabled advertising technologies.

More AI is not always better. While improved generative capability enhances advertising effectiveness in isolation, in competitive markets it can intensify rivalry and reduce profits. Firms should evaluate AI investments not only in terms of internal efficiency gains but also in terms of their strategic spillovers.

Identification and generation capabilities are distinct strategic levers. Investments in consumer message-preference identification and in generative personalization have fundamentally different competitive consequences. Identification improvements may soften competition by encouraging focus on core segments, whereas generative improvements may expand competitive overlap.

Asymmetric AI adoption may emerge naturally. Firms may not converge to identical AI technology levels, even if they start symmetrically. Observed heterogeneity in AI capability across competitors may reflect equilibrium strategic incentives rather than managerial inefficiency.

Targeting expansion may require pricing adjustments. When generative accuracy

improves, expanding advertising into more distant consumer segments may necessitate lower prices to convert marginal buyers. Managers should coordinate advertising technology investments with pricing strategy rather than treating them as independent decisions.

6.2. Limitations and Future Research

Our framework abstracts from several features that may warrant further investigation. First, we model a static setting with one-shot advertising and purchase decisions. In practice, firms learn over time, and AI systems update with feedback. A dynamic model incorporating learning and experimentation could shed light on the evolution of targeting strategies and investment incentives.

Second, we focus on a single-product firm (or symmetric duopoly) with exogenous baseline creative locations. Extending the model to multi-product firms, endogenous product positioning, or platform-mediated advertising environments could generate additional insights.

Third, while we capture generative distortion in reduced-form fashion, future work could explore richer representations of hallucination risk, brand consistency concerns, or consumer trust in AI-generated content. Incorporating consumer inference about AI usage may further affect equilibrium outcomes.

Finally, our model abstracts from privacy regulation and data constraints. As regulatory environments evolve, constraints on data usage may differentially affect identification and generative capabilities, altering firms' strategic incentives.

Overall, our analysis demonstrates that AI-enabled advertising technologies reshape not only how firms communicate with consumers, but also how they compete, invest, and price. By disentangling the roles of preference identification and generative personalization, we provide a framework for understanding the strategic implications of GenAI adoption in modern marketing environments.

Appendix A: proofs

Proof of Lemma 1. Starting with α_m . By the envelope theorem, $\frac{\partial \Pr[\text{click}|m_i, m_i^*]}{\partial \alpha_m} = \frac{1}{6} - \frac{1}{2}d(m_i, m_i^*)^2 = \frac{1}{6} - \frac{(1-\beta)^2}{2(1-\beta+\alpha_m)^2}d(m^F, m_i)^2$. Thus, $\frac{\partial \Pr[\text{click}|m_i, m_i^*]}{\partial \alpha_m} < 0$ iff $\frac{1-\beta+\alpha_m}{\sqrt{3}(1-\beta)} < d(m^F, m_i)$. Notice that when $(\sqrt{3}-1)(1-\beta) < \alpha_m$, the LHS is greater than 1 implying that $\frac{\partial \Pr[\text{click}|m_i, m_i^*]}{\partial \alpha_m} \geq 0$ and $\Pr[\text{click} | m_i, m_i^*]$ increases in α_m for all $m_i \in [0, 2]$. When $\alpha_m < (\sqrt{3}-1)(1-\beta)$, $\Pr[\text{click} | m_i, m_i^*]$ decreases in α_m iff $m_i \in (\frac{1-\beta+\alpha_m}{\sqrt{3}(1-\beta)}, 2 - \frac{1-\beta+\alpha_m}{\sqrt{3}(1-\beta)})$.

Turning now to β . By the envelope theorem, $\frac{\partial \Pr[\text{click}|m_i, m_i^*]}{\partial \beta} = \frac{1}{2}d(m^F, m_i^*)^2$. Thus, the probability of click increases with β if $\alpha_m > 0$ and $m_i \neq m^F$. The probability of click is independent of β when $\alpha_m = 0$ or $m_i = m^F$. ■

Proof of Proposition 1. We show that, when $\frac{1}{6} + \frac{1}{\sqrt{3}} < \frac{c}{\pi}$, there exist two threshold functions $\bar{\alpha}_m(\frac{c}{\pi})$ and $\bar{\beta}(\alpha_m, \frac{c}{\pi})$, such that the firm's optimal target strategy is as follows:

- $s^* = 0$ when $\frac{5}{6} < \frac{c}{\pi}$ and $\alpha_m \leq \bar{\alpha}_m(\frac{c}{\pi})$;
- $s^* \in (0, 1)$ when $\frac{1}{6} + \frac{1}{\sqrt{3}} < \frac{c}{\pi}$, $\bar{\alpha}_m(\frac{c}{\pi}) < \alpha_m$ and $\beta < \bar{\beta}(\alpha_m, \frac{c}{\pi})$;
- $s^* = 1$ otherwise.

In addition, we show that $\bar{\alpha}_m(\frac{c}{\pi})$ is decreasing for $\frac{c}{\pi} < \frac{5}{6}$ and increasing for $\frac{5}{6} < \frac{c}{\pi}$, while $\bar{\beta}(\alpha_m, \frac{c}{\pi})$ increases in $\frac{c}{\pi}$.

Consider the consumer with IdAI matching the base message of the firm: $d(m^F, m_i) = 0$. If the cost of sending a message is greater than the expected revenue for that consumer, the firm does not target anybody. Specifically, if $\frac{5+\alpha_m}{6}\pi < c \Leftrightarrow \alpha_m < 6\frac{c}{\pi} - 5$, then $s^* = 0$. If $\frac{c}{\pi} > 1$, $s^* = 0$ for all α_m and β . If $\frac{5}{6} < \frac{c}{\pi} \leq 1$, define $\bar{\alpha}_m(\frac{c}{\pi}) = 6\frac{c}{\pi} - 5$, which is increasing in $\frac{c}{\pi}$ and equal to 0 when $\frac{c}{\pi} = \frac{5}{6}$.

Now, consider the consumer with IdAI the furthest away from the base message of the firm: $d(m^F, m_i) = 1$. If the cost of sending a message is lower than the expected revenue for that consumer, the firm targets every consumer. In other words, if $c \leq \left(\frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)}\right)\pi \Leftrightarrow \frac{c}{\pi} \leq \frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)}$, then $s^* = 1$. Notice that $\frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)} \geq \frac{1}{6} + \frac{1}{\sqrt{3}}$ with equality when $\alpha_m = \sqrt{3} - 1$ and $\beta = 1$. If $\frac{c}{\pi} < \frac{1}{6} + \frac{1}{\sqrt{3}}$, $s^* = 1$ for all α_m and β . In the rest of the proof, we focus on $\frac{1}{6} + \frac{1}{\sqrt{3}} < \frac{c}{\pi}$. We can express $\frac{c}{\pi} \leq \frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)}$ as $\left[\frac{2\alpha_m-5}{6} + \frac{c}{\pi}\right](1-\beta) \leq \alpha_m \left(\frac{5+\alpha_m}{6} - \frac{c}{\pi}\right)$. If $\frac{2\alpha_m-5}{6} + \frac{c}{\pi} \leq 0 \Leftrightarrow \alpha_m \leq 3(\frac{5}{6} - \frac{c}{\pi})$, this condition is true. If $3(\frac{5}{6} - \frac{c}{\pi}) < \alpha_m$, this condition is equivalent to $\beta > 1 + \alpha_m - \frac{\alpha_m^2/2}{\frac{2\alpha_m-5}{6} + \frac{c}{\pi}} \Leftrightarrow \beta > \bar{\beta}(\alpha_m, \frac{c}{\pi})$, where $\bar{\beta}(\alpha_m, \frac{c}{\pi}) = \max[1 + \alpha_m - \frac{\alpha_m^2/2}{\frac{2\alpha_m-5}{6} + \frac{c}{\pi}}, 0]$. $\bar{\beta}$ is increasing in $\frac{c}{\pi}$ and, if $\frac{1}{6} + \frac{1}{\sqrt{3}} < \frac{c}{\pi} < \frac{5}{6}$, $\bar{\beta} > 0 \Leftrightarrow \alpha_m > \bar{\alpha}_m(\frac{c}{\pi}) = \frac{3c}{\pi} - \frac{3}{2} - \frac{1}{2}\sqrt{(6\frac{c}{\pi} - (1+2\sqrt{3}))(6\frac{c}{\pi} + (2\sqrt{3}-1))}$, which is decreasing in $\frac{c}{\pi}$ and equal to 0 when $\frac{c}{\pi} = \frac{5}{6}$.

Finally, if $\frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)} < \frac{c}{\pi} \leq \frac{5+\alpha_m}{6}$, the firm targets some consumers but not all.

Specifically, $0 < s^* < 1$, where s^* is solution to $\left(\frac{5+\alpha_m}{6} - \frac{\alpha_m(1-\beta)}{2(1-\beta+\alpha_m)}(s^*)^2\right)\pi = c \Leftrightarrow s^* = \sqrt{\frac{2(1-\beta+\alpha_m)}{\alpha_m(1-\beta)}\left(\frac{5+\alpha_m}{6} - \frac{c}{\pi}\right)}$. ■

Proof of Proposition 2. When $s^* \in (0, 1)$, recall that $(s^*)^2 = \frac{2(1-\beta+\alpha_m)}{\alpha_m(1-\beta)}\left(\frac{5+\alpha_m}{6} - \frac{c}{\pi}\right)$. Hence, s^* is decreasing in α_m iff $\alpha_m^2 < 6(1-\beta)\left(\frac{5}{6} - \frac{c}{\pi}\right) \Leftrightarrow \left(\frac{c}{\pi} < \frac{5}{6} \text{ and } \alpha_m < \sqrt{6(1-\beta)\left(\frac{5}{6} - \frac{c}{\pi}\right)}\right)$. Conversely, s^* is increasing in α_m iff $\alpha_m^2 > 6(1-\beta)\left(\frac{5}{6} - \frac{c}{\pi}\right) \Leftrightarrow \frac{c}{\pi} \geq \frac{5}{6}$ or $\left(\frac{c}{\pi} < \frac{5}{6} \text{ and } \alpha_m > \sqrt{6(1-\beta)\left(\frac{5}{6} - \frac{c}{\pi}\right)}\right)$. In addition, s^* is increasing in β because $\frac{1-\beta+\alpha_m}{1-\beta} = 1 + \frac{\alpha_m}{1-\beta}$ is increasing in β . ■

Proof of Lemma 2. We prove the result for consumers in locations $m_i \in [0, 1]$. The result holds by symmetry for consumers in locations $m_i \in [1, 2]$. We do this by proving that, in equilibrium, the probability that a consumer i in location $m_i \in [0, 1]$, who is presented with ads of both firms, clicks on Firm 0's ad

- (i) Decreases in m_i
- (ii) Increases in α_m iff $m_i < 1/2$
- (iii) Decreases in β iff $m_i < 1/2$

Recall from Section 4.1 that the optimal message $m_{j,i}^*$ is the same as m_i^* in the monopoly case, with m_j^F replacing m^F of the monopoly. That is,

$$m_{0,i}^* = \frac{\alpha_m m_i + (1 - \alpha_m)/2}{2 - \beta}; \quad m_{1,i}^* = 1 - \frac{\alpha_m(1 - m_i) + (1 - \alpha_m)/2}{2 - \beta}$$

Putting this together we get that in equilibrium the probability that consumer i with observed ad preferences m_i clicks on the ad from Firm 0 is

$$F^*(m_i) = \frac{1}{2} + \frac{\alpha_m(1 - \beta)\left(\frac{1}{2} - m_i\right)}{2(2 - \beta)},$$

and the consumer clicks firm 1's ad with the remaining probability.

To complete the proof it is sufficient to note that,

$$\frac{\partial F^*(m_i)}{\partial \alpha_m} = \frac{(1 - \beta)\left(\frac{1}{2} - m_i\right)}{2(2 - \beta)}, \quad (8)$$

$$\frac{\partial F^*(m_i)}{\partial \beta} = -\frac{\alpha_m\left(\frac{1}{2} - m_i\right)}{2(2 - \beta)^2}, \quad (9)$$

$$\frac{\partial F^*(m_i)}{\partial(m_i)} = -\frac{\alpha_m(1 - \beta)}{2(2 - \beta)}. \quad (10)$$

■

Proof of Lemma 3. We prove the result for consumers in locations $m_i \in [0, 1]$. The result holds by symmetry for consumers in locations $m_i \in [1, 2]$.

In equilibrium, the choice of s_0 does not depend on s_1 . To see why, note that as long as there is overlap ($s_0 + s_1 > 1$), the probability of any consumer with identified preference location m_i in the neighborhood of s_0 clicks Firm 1's ad is independent of s_1 . Similarly, the choice of s_1 does not depend on s_0 . Symmetry with respect to the profit functions of the two firms implies that $s_0^* = s_1^*$. We now prove the lemma by proving the result for s_0 . The proof for s_1 is symmetric.

Suppose that a consumer i , who is presented with ads of both firms, has identified preference location that is near to Firm 1 ($m_i \in (1/2, 1]$). A direct implication of Part 2 of Lemma 2 is that an improvement in the IdAI capabilities (α_m) leads to a decrease in the probability that this consumer clicks on Firm 0's ad. Similarly, a direct implication of Part 3 of Lemma 2 is that an improvement in the GenAI capabilities (β) leads to an increase in the probability that the consumer clicks/notices Firm 0's ad.

Because firms are symmetric we have that $s_0^* = s_1^*$ and we also know that $s_0^* + s_1^* > 1$ because we focus on the case that an overlap in advertising exists in equilibrium. So $d_0^* > 1/2$. As a result, the probability of a click from Firm 0's marginal consumer decreases in α_m (which leads Firm 0 to weakly reduce s_0^* – strictly if $s_0^* \neq 1$) and increases in β (which leads firm 0 to weakly increase s_0^* – strictly if $s_0^* \neq 1$). ■

Proof of Proposition 3. Lemma 3 concludes that there exists some s^* such that $s_0^* = s_1^* = s^*$. Moreover, the comparative statics for s^* derived in Lemma 3 do not affect aggregate sales (which remain = 1), and equilibrium symmetry implies that such changes do not affect either firm's revenues. On the other hand, changes in s^* directly affect costs of both firms, leading to the change in profits captured by the proposition. ■

Proof of Proposition 4. Denote by $s_j^*|_{I_0, I_1}$ the equilibrium advertising segment size of Firm j , where $I_j \in \{y, n\}$ captures if Firm j invested in GenAI.

If both firms do not invest in GenAI, then they are symmetric and we have that each firm is clicked on by half of the consumers and has a profit of $\frac{1}{2}\pi - cs_j^*|_{n,n}$.

If both firms invest in GenAI, then firms are symmetric and we have that each firm is clicked on by half of the consumers and has a profit of $\frac{1}{2}\pi - cs_j^*|_{y,y} - C$. We can say more: if $\beta_0 = \beta_1 = 1$ then the location of firm j doesn't affect V_j . Therefore, any consumer i , who is advertised to by both firms, clicks on Firm 0's ad with probability 1/2. As a result, in equilibrium, either there is no overlap between the firms' advertising intervals, or there is full overlap with $s_j^*|_{y,y} = 1$ for both

firms. Because we focus on the case that there is overlap, we have that $s_j^*|_{y,y} = 1$ and the profit of each firm is $\frac{1}{2}\pi - c - C$.

Finally, if Firm 1 doesn't invest in GenAI and Firm 0 does (the reverse case is symmetric): Firm 0 has higher return to advertising in every location than it had in the case that both firms invested, and so $s_0^*|_{y,n} = 1$. Also, for similar reasons $s_1^*|_{y,n} \leq s_1^*|_{n,n}$. In this case Firm 0's ad is clicked on by more than 1/2 of the consumers and Firm 1's ad by fewer than 1/2 of the consumers. The combined profits of both firms is $\pi - c - cs_1^*|_{y,n} - C$

Few observations:

- (i) Moving from no investment to both firms investing, both firms lose because they pay C and both weakly increase s_j^* , whereas their revenues do not change.
- (ii) When a firm moves from not investing to investing it pays C (regardless of whether the other firm invested) and it gains in sales the same regardless of whether the other firm invested (because the sum of clicks always equals 1).
- (iii) When a firm moves from not investing to investing it increases its s_j^* by
 - (i) if the other firm invests: $s_1^*|_{y,y} - s_1^*|_{y,n} = 1 - s_1^*|_{y,n}$
 - (ii) if the other firm does not invest: $s_0^*|_{y,n} - s_0^*|_{n,n} = 1 - s_1^*|_{n,n}$
- (iv) Because $s_1^*|_{y,n} \leq s_1^*|_{n,n}$ we have that s_j^* increases by more if the other firm invests, and because the increase in revenues is the same, the added profit is smaller. That is, a firm has larger incentive to invest in GenAI if the other firm does not invest. This establishes Part 1 of the proposition. Holding all other parameters fixed Part 2 directly follows.

■

Proof of Proposition 5. Forthcoming, available from the authors. ■

Proof of Proposition 6. Forthcoming, available from the authors. ■

Appendix B: Heterogeneous GenAI Capabilities

Messaging: competing on consumer i

Because firm j chooses its message to consumer i by maximizing only V_j , it bases it only on β_j . Therefore, we can directly say that for any $m_i^c \in [0, 1]$:

$$m_{0,i}^* = \frac{\alpha_m m_i + (1 - \alpha_m)/2}{2 - \beta_0}; \quad m_{1,i}^* = 1 - \frac{\alpha_m(1 - m_i) + (1 - \alpha_m)/2}{2 - \beta_1}$$

Substitute into the equilibrium probability that a consumer i , who is displayed ads by both firms, clicks/notices firm 0's ad we get that

$$F^*(y_j^{DB}) = \frac{\alpha_m}{2} [1 + V_0(m_{0,i}^*, m_i) - V_1(m_{1,i}^*, m_i)] + \frac{1 - \alpha_m}{2} \int_0^1 [1 + V_0(m_{0,i}^*, m) - V_1(m_{1,i}^*, m)] dm \quad (11)$$

$$= \frac{1}{2} + \frac{\alpha_m}{2} [V_0(m_{0,i}^*, m_i) - V_1(m_{1,i}^*, m_i)] + \frac{1 - \alpha_m}{2} \int_0^1 [V_0(m_{0,i}^*, m) - V_1(m_{1,i}^*, m)] dm \quad (12)$$

Proposition 7 *In equilibrium, the probability that a consumer i , who is presented with ads of both firms, clicks on firm j 's ad:*

- (i) *decreases in the consumer's identified distance from firm j (i.e. $d(m_j^F, m_i)$)*
- (ii) *increases in firm j 's GenAI capabilities and decreases in the the GenAI capabilities of its competitor.*

Proof.

We prove the result for consumers in locations $m_i \in [0, 1]$. The result holds by symmetry for consumers in locations $m_i \in [1, 2]$.

Part 1: we prove that, in equilibrium, the probability that a consumer i , who is presented with ads of both firms, clicks on firm 0's ad, i.e., $F^*(m_i)$ decreases in m_i . Using the envelope theorem, we have that $\frac{dF^*(m_i)}{dm_i} = \frac{\partial F^*(m_i)}{\partial m_i} = \frac{\alpha_m}{2} (m_{0,i}^* - m_{1,i}^*)$. Hence, to prove the aforementioned claim it is sufficient to note that for any $m_i \in [0, 1]$ we have that $m_{0,i}^* < m_{1,i}^*$. By symmetry, the claim holds for firm 1.

Part 2: Consider an increase in β_0 . Holding fixed $m_{0,i}^*, m_{1,i}^*$, this leads to an increase in V_0 and therefore an increase $F^*(m_i)$. Suppose by contradiction that, in equilibrium (i.e., allowing $m_{0,i}^*, m_{1,i}^*$ to change), we have that $F^*(m_i)$ weakly decreases, and recall that β_0 does not enter firm

1's optimization. Then this implies that $m_{0,i}^*$ is not chosen optimally, contradicting that this is an equilibrium. This completes the proof that the probability that a consumer i , who is presented with ads of both firms, clicks on firm 0's ad increases in β_0 . This also proves that the probability that a consumer i , who is presented with ads of both firms, clicks on firm 1's ad decreases in β_0 . By symmetry these two claims hold substituting 0 and 1 in the firms' and β 's subscripts. ■

Targeting: who to display ads to

Proposition 8 *An increase in β_j increases the advertising interval of firm j and decreases the advertising interval of its competitor.*

Proof. We prove the result for consumers in locations $m_i \in [0, 1]$. The result holds by symmetry for consumers in locations $m_i \in [1, 2]$.

We prove the proposition by showing that:

- (i) s_j^* is increasing in β_j
- (ii) s_0^* is decreasing in β_1 and s_1^* is decreasing in β_0

Part 1: With some abuse of notation let $m_j^*(m)$ be the equilibrium message displayed by firm j to a consumer with identified preference m and note that $\frac{\partial^2 \Pi_0}{\partial s_0 \partial \beta_0} = \frac{\alpha_m}{2} \frac{\partial V_0(m_0^*(s_0), s_0)}{\partial \beta_0} + \frac{1-\alpha_m}{2} \int_0^1 \frac{\partial V_0(m_0^*(s_0), m)}{\partial \beta_0} dm$ which is positive because $\frac{\partial V_0(m_0^*(s_0), m)}{\partial \beta_0} > 0$ for all m .

Part 2: Similarly to 1, note that $\frac{\partial^2 \Pi_0}{\partial s_0 \partial \beta_1} = -\frac{\alpha_m}{2} \frac{\partial V_1(m_1^*(s_0), s_0)}{\partial \beta_1} - \frac{1-\alpha_m}{2} \int_0^1 \frac{\partial V_1(m_1^*(s_0), m)}{\partial \beta_1} dm$ which is negative because $\frac{\partial V_1(m_1^*(s_0), m)}{\partial \beta_1} > 0$ for all m . ■

References

- K. Bagwell. The economic ananlysis of advertising. In M. Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3, pages 1701–1844. Elsevier, 2007. doi: 10.1016/S1573-448X(06)03028-7.
- D. Bergemann and A. Bonatti. Targeting in advertising markets: Implications for offline versus online media. *RAND Journal of Economics*, 42(3):417–443, 2011.
- Y. Chen, C. Narasimhan, and Z. J. Zhang. Individual marketing with imperfect targetability. *Marketing Science*, 20(1):23–41, 2001.
- R. Dolan, J. Avery, and C. Chammas. Meta: Digital marketing and articial intelligence (ai) at facebook and instagram. *Harvard Business Publishing case*, 2025. HBS case number 526-031.
- M. Elias. Video ads trends in 2025: Next in AI, personalization and data. *Forbes*, March, 14 2025. <https://www.forbes.com/councils/forbesbusinesscouncil/2025/03/04/video-ads-trends-in-2025-next-in-ai-personalization-and-data/>, accessed October 8, 2025.
- I. P. Fainmesser, A. Galeotti, and R. Momot. Consumer profiling via information design. *Available at SSRN 4655468*, 2023.
- P. M. Gardete and Y. Bart. Tailored cheap talk: The effects of privacy policy on ad content and market outcomes. *Marketing Science*, 37(5):685–853, 2018.
- Y. Gil and P. Raymond. The AI index 2025 annual report. *AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA*, 2025.
- D. Grewal, C. B. Santornino, T. H. Davenport, and A. Guha. How should gen ai fit into your marketing strategy? *Harvard Business Review*, March, 24 2025.
- G. M. Grossman and C. Shapiro. Informative advertising with differentiated products. *Review of Economic Studies*, 51(1):63–81, 1984.
- J. R. Hauser, G. L. Urban, G. Liberali, and M. Braun. Website morphing. *Marketing Science*, 28 (2):202–223, 2009.
- G. Iyer, D. Soberman, and J. M. Villas-Boas. The targeting of advertitising. *Marketing Science*, 24 (3):461–476, 2005.

- A. Kapoor and M. Kumar. Generative AI and personalized video advertisements. *Marketing Science*, 44(4):733–747, 2025.
- S. Matz, J. D. Teeny, S. S. Vaid, H. Peters, G. Harari, and M. Cerf. The potential of generative AI for personalized persuasion at scale. *Scientific Reports*, 14(4692), 2024.
- Y. Ngugi. 50+ AI marketing statistics in 2025: AI marketing trends insights. *seo.com*, 2025. <https://www.seo.com/ai/marketing-statistics/>, accessed October 8, 2025.
- R. Pieters, M. Wedel, and R. Batra. The stopping power of advertising: Measures and effects of visual complexity. *Journal of Marketing*, 74:48–60, September 2010.
- J. Shin and J. Yu. Targeted advertising and consumer inference. *Marketing Science*, 40(5):900–922, 2021.
- E. Stein, K. Robinson, A. Wolfer, G. Almeida, and W. Huang. Unlocking the next frontier of personalized marketing. *McKinsey Quarterly*, 2025. January 30.
- Y. Sun. The exact law of large numbers via fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1):31–69, 2006.
- Y. Sun and Y. Zhang. Individual risk and lebesgue extension without aggregate uncertainty. *Journal of Economic Theory*, 144(1):432–443, 2009.
- K. Zhang and Z. Katona. Contextual advertising. *Marketing Science*, 31(6):980–994, 2012.