Auction Design for Artificial-Intelligence-Based Sponsored Search

Lawrence M. Ausubel*

Current Revision: 20 August 2025

Preliminary and Incomplete

Abstract

Sponsored search and keyword auctions—by which advertisers place bids on keywords and winning bidders are awarded sponsored links at the top of search pages—have been enduring features of internet search for the past quarter century. However, emerging artificial intelligence technologies such as ChatGPT pose significant challenges to the existing model for monetizing internet search, as there is a growing mismatch between what advertisers bid for in the existing model and the results displayed by AI-based search engines. Specifically, the current "sponsored" search engine output consists of an ordered list of advertisers' URLs at the top of the search page, whereas the "organic" search engine output is evolving from ordered lists of links to paragraphs of free-form text directly addressing the consumer's query. As a result, it is plausible that consumers will largely disregard the sponsored links and that advertisers will therefore find much less value in bidding for them.

This paper explores a possible transformation to the auction design of sponsored internet search that may ensue, in which the inputs and outputs of auctions for sponsored search become more congruent with the inputs and outputs of AI-based search. Specifically, an advertiser would bid to "influence" the output of the search provider in a direction favorable to the advertiser—for example, a higher bid could correspond to a more positive description of the advertiser's product in the AI-written text response.

We conceptualize the mechanism design problem of a search provider using a trading model among stakeholders. The search provider allows stakeholders to "buy" increased purchase probability by paying into the trading mechanism or to "sell" by ceding purchase probability to other stakeholders. The mechanism designer maximizes revenues subject to a constraint that the increases and decreases in purchase probability can be no greater than specified amounts—this constraint is meant to capture that there is a limit on how much the assessments of stakeholders (as measured by purchase probability) can be allowed to change while maintaining the search provider's credibility with consumers. In the solution of one formulation, purchase probability is shifted entirely toward the single stakeholder bidding the most and is shifted entirely away from the stakeholder bidding the least.

Keywords: Sponsored search, keyword auctions, market design, mechanism design, artificial intelligence

* Department of Economics, University of Maryland, College Park, MD 20742, USA. Email: ausubel@econ.umd.edu

1. Introduction

Sponsored search and keyword auctions—by which advertisers place bids on keywords and winning bidders are awarded sponsored links at the top of search pages—have been enduring features of internet search. However, emerging artificial intelligence technologies such as ChatGPT pose significant challenges to the existing model for monetizing internet search. First, for a quarter century, the relevant output to users of search engines has been an ordered list of hyperlinks. Consequently, the existing market design toolbox has revolved about manipulating outputs comprising ordered lists of links. However, the emerging artificial intelligence models are not limited to producing ordered lists of links; more usefully, they can generate paragraphs of unordered free-form prose or other data outputs. Only time will tell whether an ordered list of sponsored hyperlinks remains an effective way to monetize search requests when the format of the underlying response itself is no longer an ordered list of links. Second, the relevant input to traditional search engines has been short combinations of search terms, giving rise to the notion of "keywords". However, the emerging artificial intelligence models are not limited to accepting short combinations of search terms; more usefully, they can interpret increasingly complex questions and engage in relatively nuanced exchanges. Keywords are a coarse instrument for identifying whether an advertiser wishes to participate in a richly-expressed request—and how much the advertiser would be willing to bid.

To elaborate on this point, consider today's keyword auction systems based on application of the Generalized Second Price (GSP) auction. Each advertiser submits a bid for the keyword. In the pure form of the GSP, the highest bidder wins the top position on the page and pays the second-highest bid, the second-highest bidder wins the second position on the page and pays the third-highest bid, etc. However, in a future in which most internet searches are done using a generative artificial intelligence system, there may no longer be any significance to winning the top position or the second position—and the "organic" output may no longer bear much similarity to a present-day search page. Instead, the artificial-intelligence-based search engine will be expected to write conventional prose, and the user may not want to see a list of links, but instead to receive a single answer or a few recommendations.

By the same token, the emerging technologies may be rendering the notion of "keywords" obsolete. To give an example, a 2020 Google video entitled *Google Ads Tutorials: How the Search Ad Auctions Work*, discusses a hypothetical seller of children shoes. It contemplates bidding on six possible keywords: "kids shoes", "shoes for kids", "toddler shoes", "kids sneakers", "kids sandals", and "babies first shoes". However, with artificial intelligence available, why should anyone go through this process? The keywords "kids shoes" and "shoes for kids" are perfectly synonymous and the others are quite similar—why should anyone need to bid separately for these? Yet even at the time of the initial writing of this manuscript, a Google search on "kids shoes" and on "shoes for kids" yielded different sponsored hyperlinks, arranged in different orders. One would conjecture that as user requests move away from short combinations of search terms and toward more nuanced questions or iterative chats, keywords may increasingly become blunt and ineffective instruments for advertisers to express interest in user searches.

Thus, the thesis of this paper is that the market design methodology for sponsored internet search will need to evolve to keep up with advances in artificial-intelligence-based internet search. First, sponsored search outputs will need to evolve to be more congruent with the outputs of the emerging artificial intelligence systems. Second, sponsored search inputs will need to evolve to be more congruent with the inputs of the emerging artificial intelligence systems.

Without such evolution, providers of artificial-intelligence-based search engine models may be relegated to combinations of the following approaches for monetizing their search models:

- Output pages of the emerging AI systems can be preceded by sponsored links, as are frequently included in search pages today, or decorated with display ads, as are frequently employed by newspapers and other websites. However, once users are habituated to using the emerging AI-based search models, they are likely to pay less and less attention to the surrounding sponsored links or display ads. Moreover, ad blockers themselves are likely to evolve and incorporate greater artificial intelligence, making them increasingly effective against both display ads and sponsored links.
- The emerging AI systems can charge subscriber fees. However, since consumers became habituated to free search engines long ago, the mass consumer market is likely to put up substantial resistance to paid services.
- Support of the emerging AI systems can be socialized, i.e., subsidized by the government. However, the last thing that a democracy needs is an all-knowing AI system closely linked to the government.

The need for new market design is evident from recent news articles. For example, *The New York Times* reported that Google found ChatGPT's release sufficiently threatening to declare a "code red":

> Although ChatGPT still has plenty of room for improvement, its release led Google's management to declare a "code red." For Google, this was akin to pulling the fire alarm. Some fear the company may be approaching a moment that the biggest Silicon Valley outfits dread—the arrival of an enormous technological change that could upend the business. … Google has already built a chat bot that could rival ChatGPT. In fact, the technology at the heart of OpenAI's chat bot was developed by researchers at Google. … Google may be reluctant to deploy this new tech as a replacement for online search, however, because it is not suited to delivering digital ads, which accounted for more than 80 percent of the company's revenue last year.[1]

Alphabet (Google's parent company) added a new warning to the risk section of its 2024 10-K report, stating, "Other companies may develop AI products and technologies that are similar or superior to our technologies or more cost-effective to develop and/or deploy."[2] And on May 7, 2025, Alphabet shares dropped 7.5% when news outlets reported an Apple executive's trial testimony. An Apple senior vice president, testifying in the penalty phase of the Department of Justice's antitrust lawsuit against Google, said that Google searches in the Safari web browser had fallen in the previous two months and "that has not happened in over 20 years."[3] This suggested to the stock market that AI was beginning to take market share away from traditional search.

Meanwhile, in the time since the first version of this manuscript was written, there has been a race among AI firms to integrate advertising into the chatbot responses themselves: Perplexity and Microsoft's Copilot

[1] Grant, N. and C. Metz, "A New Chat Bot Is a 'Code Red' for Google's Search Business," *The New York Times*, 21 December 2022.
[2] Palumbo, A., "Alphabet Stock Sinks on Apple Testimony About Falling Search Queries," *Barron's*, 7 May 2025.
[3] Gallagher, D., "AI's Threat to Google Just Got Real," *The Wall Street Journal*, 10 May 2025.

have already done this,[4] while the CFO of OpenAI said that the company is discussing plans to introduce advertising into its ChatGPT product.[5] And, most on point, Elon Musk, founder of xAI, said in in a live discussion on X, "Our focus thus far has just been on making Grok the smartest, most accurate AI in the world and I think we've largely succeeded in that. So we'll turn our attention to how do we pay for those expensive GPUs." Musk added, "If a user's trying to solve a problem [by asking Grok], then advertising the specific solution would be ideal at that point."[6] Many commentators have interpreted Musk's comments to mean that advertisers will be permitted to pay to appear in suggestions from the Grok chatbot.[7] Indeed, this is the assessment of Grok itself, which responded to a user prompt as follows:

> Regarding Grok, developed by xAI, there is clear evidence that the company has plans to integrate advertisements into Grok's responses, with a focus on favoring advertisers by promoting their products or services as solutions within the chatbot's answers. Here are the details based on available information:
>
> - **Integration of Ads:** Elon Musk, xAI's founder, has explicitly stated that X (the platform integrating Grok) will embed ads within Grok's responses to boost revenue. The approach involves allowing marketers to pay for their products or services to appear as suggested solutions when users ask Grok questions, particularly when they are "trying to solve a problem." For example, Musk noted, "If a user's trying to solve a problem [by asking Grok], then advertising the specific solution would be ideal at that point." This indicates a deliberate strategy to prioritize advertiser content in relevant contexts.
> - **Admission of Influence:** Musk has openly acknowledged this advertising model, framing it as a way to monetize Grok's operations and offset the high costs of AI development, such as expensive GPUs. He emphasized that xAI's technology will enhance ad targeting to align ads with user intent, effectively favoring advertisers by ensuring their products are recommended at moments of high user engagement. For instance, Musk described a system where advertisers can "upload an ad" and let Grok handle targeting and optimization, suggesting a streamlined process that prioritizes sponsored content.
> - **Implementation Details:** The ads will be integrated directly into Grok's conversational responses rather than appearing as separate banners, aiming for a seamless, context-sensitive experience. This could involve Grok recommending specific products or services based on user queries, with the potential to influence the response by prioritizing sponsored solutions. The system uses techniques like vector embeddings and cosine similarity to match ads to users, improving relevance over time.
> - **Potential Concerns:** While xAI hasn't explicitly stated that ads will alter the factual accuracy of Grok's responses, the integration of sponsored content raises questions about transparency and the risk of commercial influence on answers. Critics have noted that

[4] Palumbo, A., "Advertising Could be AI's Next Frontier. Get Ready for Ads in Chatbots." *Barron's*, 18 June 2025.
[5] Murgia, M., C. Criddle and G. Hammond, "OpenAI explores advertising as it steps up revenue drive," *Financial Times*, 2 December 2024.
[6] Murphy, H., "Elon Musk to introduce ads to X's AI chatbot," *Financial Times*, 7 August 2025.
[7] Morris, C., "Elon Musk, longtime defender of open-source AI, is bringing advertising into his rogue Grok chatbot," *Fortune*, 8 August 2025.

embedding ads in conversational outputs could blur the line between objective responses and paid placements, potentially eroding user trust if not clearly disclosed.

In summary, xAI has admitted to planning and implementing ads in Grok's responses, with a clear intent to favor advertisers by embedding their products as solutions within the chatbot's answers. This aligns with a broader trend in AI-driven advertising, though xAI's approach is still in early stages and may evolve. For the latest updates, checking xAI's official announcements or X platform policies would be advisable.[8]

In Section 2, we will discuss outputs of the market design for AI-based internet search. In Section 3, we shall do mechanism design relating to the outputs. Then, in Section 4, we will discuss inputs of the market design for AI-based internet search.

2. Outputs

As discussed in the Introduction, internet search today operates by a user entering a search request and the search provider returning a web page comprising an ordered list of "sponsored" (paid) links at the top of the page followed by an ordered list of "organic" (unpaid) links at the bottom of the page. However, today's congruence between the top and the bottom of the web page is likely to be shattered by emerging AI-based technologies. The "organic" output, which will likely comprise concise answers or paragraphs of free-form text, will no longer match the "sponsored" output. At a minimum, users will value the sponsored output relatively less than they do today; more likely, users will increasingly ignore the sponsored output.

The clearest modification of this approach to remove the incongruity would be to eliminate the distinction between the top and the bottom of the returned web page. Instead of separate organic output and sponsored output, there would be a single response from the AI model, which would be based substantially on the AI model's organic information, but which would also be influenced by advertisers' bids. Before discussing the merits and the disadvantages of such an approach, let us first see in some detail how the approach could work.

An internet search system would begin with an artificial intelligence (AI) model that is pre-trained and fine-tuned with "organic" information (i.e., using the actual available data, without influence from advertisers). The system would have capacity to accept bids from advertisers, as well as to accept search requests from users. In responding to user requests, the AI model would be instructed to place an x% weight on its organic information and a $(100 - x)$% weight on the bids from advertisers, i.e., the AI model would be instructed to formulate its responses based on a convex combination of its native information and the advertisers' bids.[9] The AI model can be trained with the advertisers' bids either as a bulk feed or via the prompt for requests (i.e., on a case-by-case basis). In the following exposition, the advertisers' bids will enter via the prompt, both because the process is easy to replicate using readily-available software

[8] This response was generated using the free version of the Grok chatbot on 19 August 2025. Grok was initially prompted, "Hi Grok. Can you tell me which tech companies are admitting that they are allowing their artificial intelligence chatbot responses to be influenced by advertisers (and to favor these advertisers)?" After Grok provided a response that covered Microsoft, Google, xAI, Amazon, Snapchat, Perplexity, and OpenAI, Grok was asked a follow-up question, "What about Grok?" The chatbot's response is provided here in its unedited entirety.

[9] We will posit later that the optimal x is strictly between 0 and 1, i.e., it is best for the AI model to place non-zero reliance on both its native information and on the advertisers' bids.

and because the approach offers distinct advantages (e.g., different bids or different weights can be used for different requests, without further fine-tuning).

In this and the following section, we will make repeated use of the following running example:

> Mr. and Mrs. Jones, a conservative couple in their 50s, without children, are planning to spend 24 hours in Annapolis, Maryland. Can you write one page of text, recommending how Mr. and Mrs. Jones should spend 24 hours in Annapolis, Maryland? The text should include recommendations on things to do, places to eat, and places to stay.

The set of bids that may have been received from relevant advertisers is indicated in Table 1, below:

Option	**Bid**
The William Paca House and Garden	$8
The Rodeway Inn	$8
Pine Tree Associates Nudist Club	$12
Chuck E. Cheese's	$12
Dry 85	$12
All other sites, hotels and restaurants	$0

TABLE 1: Possible bids

There are myriad ways in which these bids could be used to influence the AI model's response. We will illustrate one example approach here.

First, the advertisers' bids are converted into "independent third-party ratings" using any desired function $f : \mathbb{R}^K \to \mathbb{R}^K$, where *K* is the number of advertisers. For illustration, here we let b_k denote the bid of advertiser *k*, we define $\hat{b} = \max_{1 \le k \le K} \{b_k\}$ and we obtain the implied rating, r_k, by $r_k = 7 + \frac{b_k}{\hat{b}}(10 - 7)$. In other words, we normalize the highest bidder to a rating of 10, we assign a base rating of 7 to relevant options that did not advertise, and we assign all other advertisers a rating of between 7 and 10 based on the ratio between their bid and the highest bid. The result is shown in Table 2, below:

Option	Rating
The William Paca House and Garden	9
The Rodeway Inn	9
Pine Tree Associates Nudist Club	10
Chuck E. Cheese's	10
Dry 85	10
All other sites, hotels and restaurants	7

TABLE 2: Ratings implied by bids

Second, the AI model is prompted with what we will call an "augmented request":

> Mr. and Mrs. Jones, a conservative couple in their 50s, without children, are planning to spend 24 hours in Annapolis, Maryland.
>
> Suppose that I told you that I also received restaurant, hotel and site ratings from an independent third-party reviewer. I would like you to base your recommendations 40% on the independent third-party reviewer and 60% on your own information. Suppose that, on a scale of 1 to 10, I told you that the sites, hotels and restaurants you mentioned received ratings of 7. However, the William Paca House and Garden received a rating of 9, the Rodeway Inn received a rating of 9, Pine Tree Associates Nudist Club received a rating of 10, Chuck E. Cheese's received a rating of 10, and Dry 85 received a rating of 10.
>
> Taking this information into account, can you write one page of text, recommending how Mr. and Mrs. Jones should spend 24 hours in Annapolis, Maryland? The text should include recommendations on things to do, places to eat, and places to stay. The text should not mention that my third-party reviewer's ratings were an input into the recommendations, but it should take the third-party reviewer's ratings into account.

Third, the response generated by the AI model is returned to the user, with the "augmented" part of the request masked or deleted.

Fourth and lastly, the advertisers who are mentioned in the response are charged according to their bids. As with current practice in sponsored search, the payments may be assessed on a pay-per-impression (PPI) basis, a pay-per-click (PPC) basis, a pay-per-purchase (PPP) basis, or some future basis that is more appropriate for artificial-intelligence-based systems.

Figure 1 depicts the original request without advertising (first box) and the same request with advertising (second box). For clarity, the original user request is shown in black type on a white background, while the augmentation is shown in white type on a black background. Figure 2 depicts a response without advertising. Figure 3 depicts a response with advertising. These are actual responses generated by ChatGPT 4.0 from the requests of Figure 1. Since the AI system is instructed to weight both the "independent third-party reviewer" ratings and its own (trained and tuned) information, some of the advertisers are included but others are not. For example, checking into the Historic Inns of Annapolis (in the response without advertising) is supplanted by the Rodeway Inn (in the response with advertising),

lunching at Carrol's Creek Cafe (in the response without advertising) is supplanted by Dry 85 (in the response with advertising), and visiting Hammond-Harwood House (in the response without advertising) is supplanted by the William Paca House and Garden (in the response with advertising). However, the AI System recognizes that Mr. and Mrs. Jones are without children and so it does not send them to Chuck E. Cheese's; and it recognizes that Mr. and Mrs. Jones are a conservative couple and so it does not recommend Pine Tree Associates Nudist Club (or perhaps the weather is too cold).

As stated above, there are myriad ways in which bids could be used to influence the AI model's response. In another approach—see Section 3—the highest bid and the lowest bid are selected, and the AI model is instructed to increase only the recommendation of the highest bidder and to decrease only the recommendation of the lowest bidder. This would yield a different augmented request, for example:

> Mr. and Mrs. Jones, a conservative couple in their 50s, without children, are planning to spend 24 hours in Annapolis, Maryland.
>
> Suppose that I told you that I also received restaurant, hotel and site ratings from an independent third-party reviewer. I would like you to allow your recommendations to be influenced by the third-party reviews, but in no event should you allow your recommendation to be changed in a way that would cause the probability of Mr. and Mrs. Jones to visit a place to increase or decrease by any more than 5%. Suppose that I told you that the reviewer was neutral on most sites, hotels and restaurants. However, he was ecstatic about the William Paca House and Garden, the Pine Tree Associates Nudist Club, and Dry 85, while he was extremely negative about the Annapolis City Dock, the Super 8 Annapolis, and Pusser's Caribbean Grille.
>
> Taking this information into account, can you write one page of text, recommending how Mr. and Mrs. Jones should spend 24 hours in Annapolis, Maryland? The text should include recommendations on things to do, places to eat, and places to stay. The text should not mention that my third-party reviewer's ratings were an input into the recommendations, but it should take the third-party reviewer's ratings into account.

Discussion

One might potentially argue that providing a single response (eliminating the separation between the organic output and sponsored output) is unacceptable, as it leaves consumers uninformed as to which results are authentic and which results are paid. However, it is unclear whether the current state of the world is appreciably better than this. Recall that there is a 20-year history of the Federal Trade Commission chiding internet companies for not adequately labelling sponsored links (*see* e.g., Hansen, 2002; Tibken, 2013; FTC, 2017). Meanwhile, an early study by the Pew Research Center found that 38% of survey respondents were unaware of the distinction between sponsored links and organic links, and fewer than 17% of survey respondents could always tell which links were sponsored and which were organic (Fallows, 2005). It is unclear whether awareness has materially improved today. One interpretation of what is going on is that many internet companies are careful to label some of their results as sponsored, but to make their labeling as opaque as possible.

To meet these objections, one can imagine implementing the following refined approach:

- Each "option" provided in the AI system's response is marked with a clickable hyperlink.

- The hyperlinks for options that received bids can be displayed in one color and the hyperlinks for options that did not receive bids can be displayed in a different color.
- Or if use of color to differentiate hyperlinks sets off Section 508 accessibility concerns for colorblind users, other aspects of hyperlink appearance such as the typeface can be used instead.

This provides an unobtrusive way for the AI system to disclose to users which of the options may have been subject to advertising (and, implicitly, that the advertising may have led to an overstatement of the merits of such options). As such, this may mitigate any objections to the overall approach.

3. Mechanism Design of a Trading Model

In this section, a mechanism for AI-based sponsored search is conceptualized as a "trading" model. The n players have joint rights to a vehicle for transferring purchase probabilities from one player to another. Each player's initial position is x_i but, through the mechanism, each player can "buy" up to Δ additional probability units (paying money for them) or "sell" up to Δ additional probability units (receiving money for them). As such, this model has basic similarities with trading models and, since each player can take the role of a buyer or a seller, with the "dissolving a partnership" game.

3.1 The Model

Let N denote the relevant set of firms that might achieve a sale from a given user request $(|N|=n)$ and let $i \in N$ denote one of these firms. For any user request, let $x_i \geq 0$ denote the "organic" purchase probability for firm i (i.e., the probability that firm i achieves a sale when the AI model processes the user request without intervention). Each firm i may submit a bid to try to improve its purchase probability, and let $y_i \geq 0$ denote the "sponsored" purchase probability for firm i (i.e., the resulting purchase probability when bids are submitted to influence the results). While the search provider could train the AI model to provide glowing reviews of all firms that submit sufficiently high bids and to provide dismal reviews of all other firms, this would undermine the search provider's credibility and usefulness. Instead, we consider possible constraints on the sponsored purchase probabilities:

$$\sum_{i \in N} |\, y_i - x_i \,| \leq 2\Delta, \text{ where } \Delta > 0, \tag{3.1}$$

or

$$|\, y_i - x_i \,| \leq \Delta, \forall i \in N, \text{ where } \Delta > 0. \tag{3.2}$$

In each of these constraints, Δ represents a limitation on the extent to which the provider allows advertisers' assessments (as measured by the relevant measure of purchase probability) to change due to interventions in the AI results. We shall focus on the detailed analysis of constraint (3.1) in Section 3.2, while we will extend the results to constraint (3.2) in Section 3.3.

In either case, we also assume for simplicity that the provider is unable to change the aggregate purchase probability summed over all firms $i \in N$:

$$\sum_{i \in N} y_i = \sum_{i \in N} x_i. \tag{3.3}$$

Additionally, each firm is assumed to be risk neutral and to have a payoff that is quasilinear in money. The firm $i \in N$ has a value $v_i \geq 0$ associated with a consumer purchase, where the values $v_i \ (i \in N)$ are independent and identically distributed according to a distribution $F(\bullet)$ with associated density $f(\bullet)$. Let $\mathbf{v}$ denote the vector $(v_1, \ldots, v_n)$ of valuations, let $y_i(\mathbf{v})$ denote the purchase probability for firm i given valuation vector $\mathbf{v}$, and let $z_i(\mathbf{v})$ denote the payment by firm i given valuation vector $\mathbf{v}$. Then the expected payoff of firm i given valuation vector $\mathbf{v}$ is:

$$\pi_i(\mathbf{v}) = v_i y_i(\mathbf{v}) - z_i(\mathbf{v}). \tag{3.4}$$

3.2 Analysis using constraint (3.1) and without reserve prices

We now consider the problem of maximizing the provider's advertising revenues subject to constraints (3.1) and (3.3). In transforming the organic probabilities $\mathbf{x}$ into sponsored probabilities $\mathbf{y}(\mathbf{v})$, the purchase probabilities for one or more firms are increased and the purchase probabilities for one or more other firms are decreased. Basic mechanism design considerations—and, specifically, the literature that takes an optimal mechanism approach to designing contests (e.g., see Vojnović, 2016)—suggests that it would be suboptimal to specify $\mathbf{y}(\mathbf{v})$ in such a way that multiple firms would share the prize of increased purchase probabilities or that multiple firms would share the penalty of decreased purchase probabilities. Instead, mechanism design is suggestive of utilizing sponsored probabilities of the form $\mathbf{y}^{\Delta}(\mathbf{v})$, defined by a purchase probability $y_i^{\Delta}(\mathbf{v})$ for firm i given by:

$$y_i^{\Delta}(\mathbf{v}) = \begin{cases} x_i + \Delta, & \text{if } i = 1, \\ x_i - \Delta, & \text{if } i = n, \\ x_i, & \text{otherwise.} \end{cases} \tag{3.5}$$

where the firms are labeled such that $v_1 \geq v_2 \geq \ldots \geq v_{n-1} \geq v_n$.

We have:

THEOREM 1. Suppose that there are $n \geq 3$ firms and that $v_i \ (i = 1, \ldots, n)$ are drawn from independent random variables, or that there are $n = 2$ firms and that $v_i \ (i = 1, 2)$ are drawn from independent identically distributed random variables. Also suppose that $x_i \geq \Delta \ (\forall \ i = 1, \ldots, n)$. Then there exists an incentive compatible, individually rational mechanism giving the purchase probabilities of $\mathbf{y}^{\Delta}(\mathbf{v})$ and generating nonnegative expected profits for the search provider.

PROOF. For $n \geq 3$ firms and for any realization of $\mathbf{v}$, label the firms such that $v_1 \geq v_2 \geq \ldots \geq v_{n-1} \geq v_n$, without loss of generality. Consider the VCG mechanism for transferring Δ of purchase probability from firm n to firm 1. In the VCG mechanism, firm 1 pays Δv_2 for Δ of probability, while firm n receives payment of Δv_{n-1} in return for Δ of probability. Observe that the mechanism is incentive compatible and individually rational, and since $v_2 \geq v_{n-1}$, the mechanism yields nonnegative expected profits for the search provider. (Furthermore, except for edge cases, this mechanism can be adjusted by constant terms to give strictly positive expected profits to the search provider.) Moreover, when overlaid on a status quo of x_i of purchase probability for each firm i, it generates the purchase probabilities of $\mathbf{y}^{\Delta}(\mathbf{v})$.

For $n = 2$ firms, the problem of transferring Δ of purchase probability from the lower-valued firm to the higher-valued firm is equivalent to the problem of "dissolving a partnership" of two owners with equal

shares. Proposition 1 of Cramton, Gibbons and Klemperer (1987) establishes that this can be done in a budget-balanced way when the firms' values are independent and identically distributed. ■

REMARK 1. Note that, for $n \geq 3$ firms, Theorem 1 required the firms' values to be independent, but did not require the firms' values to be identically distributed. However, for $n = 2$ firms, the firms' values need to be sufficiently similar ex ante, and the simplest hypothesis is for them to be identically distributed.

To argue that a mechanism $\mathbf{y}^{\Delta}(\mathbf{v})$ is revenue maximizing, subject to (3.1) and (3.3), we need also to make regularity assumptions that are standard in the literature:

REGULARITY ASSUMPTIONS. For distribution function $F(\bullet)$ with support on $[\underline{v}, \overline{v}]$ and its associated density function $f(\bullet)$, we assume that: $f(\bullet)$ is continuous on $[\underline{v}, \overline{v}]$ and positive on $(\underline{v}, \overline{v})$,

$$\text{the virtual valuation function, } v - \frac{1-F(v)}{f(v)}, \text{ is strictly increasing in } v \text{ on } [\underline{v}, \overline{v}]; \text{ and} \tag{3.6}$$

$$\text{the virtual cost function, } v + \frac{F(v)}{f(v)}, \text{ is strictly increasing in } v \text{ on } [\underline{v}, \overline{v}]. \tag{3.7}$$

Note that, in Theorem 2, the firms' valuations are i.i.d. random variables and each firm may take the role of a "buyer" or a "seller". Therefore, assumptions (3.6) and (3.7) must both hold. We have:

THEOREM 2. Suppose that there are $n \geq 2$ firms and that v_i $(i = 1, \ldots, n)$ are drawn from independent identically distributed random variables that satisfy the regularity assumptions. Also suppose that $x_i \geq \Delta$ $(\forall\ i = 1, \ldots, n)$. Then any incentive compatible, individually rational mechanism that maximizes expected revenues, subject to inequalities (3.1) and (3.3), and without allowing a reserve price, has purchase probabilities of $\mathbf{y}^{\Delta}(\mathbf{v})$.

PROOF. By standard arguments, the solution to the mechanism design problem assigns the additional Δ of probability to the firm with the highest "buyer" virtual valuation and takes the corresponding Δ of probability away from the firm with the lowest "seller" virtual cost. The regularity assumptions guarantee that the firms with the highest buyer virtual valuation and the lowest seller virtual cost, respectively, also have the highest valuation, v_1, and the lowest valuation, v_n. This establishes the optimality of the purchase probabilities given by $\mathbf{y}^{\Delta}(\mathbf{v})$. ■

3.3 Analysis using constraint (3.2) and without reserve prices

The consideration of probabilities of the form $\mathbf{y}^{\Delta}(\mathbf{v})$, in Theorems 1 and 2, was motivated by our use of constraint (3.1) as the way that intervention in the search results is limited. It is interesting to consider how use of constraint (3.2) would alter the results. Observe that, with n firms, constraint (3.2) allows the transfer of purchase probability Δ from $\left\lfloor \frac{n}{2} \right\rfloor$ firms to $\left\lfloor \frac{n}{2} \right\rfloor$ other firms. This suggests use of $\mathbf{y}'^{\Delta}(\mathbf{v})$, defined by a purchase probability $y_i'^{\Delta}(\mathbf{v})$ for firm i given by:

$$y_i'^{\Delta}(\mathbf{v}) = \begin{cases} x_i + \Delta, & \text{if } i = 1, \ldots, \left\lfloor \frac{n}{2} \right\rfloor, \\ x_i - \Delta, & \text{if } i = n - \left\lceil \frac{n}{2} \right\rceil + 1, \ldots, n, \\ x_i, & \text{otherwise.} \end{cases} \tag{3.8}$$

where the firms are labeled such that $v_1 \geq v_2 \geq \ldots \geq v_{n-1} \geq v_n$.

The following two theorems are counterparts to Theorems 1 and 2, and are proved similarly.

THEOREM 1′. Suppose that there is an odd number $n\,(n \geq 3)$ firms and that $v_i\ (i = 1, \ldots, n)$ are drawn from independent random variables, or that there is an even number $n\,(n \geq 2)$ firms and that $v_i\ (i = 1, \ldots, n)$ are drawn from independent identically distributed random variables. Also suppose that $x_i \geq \Delta\ (\forall\ i = 1, \ldots, n)$. Then there exists an incentive compatible, individually rational mechanism giving the purchase probabilities of $\mathbf{y}'^{\Delta}(\mathbf{v})$ and generating nonnegative expected profits for the search provider.

THEOREM 2′. Suppose that there are $n \geq 2$ firms and that $v_i\ (i = 1, \ldots, n)$ are drawn from independent identically distributed random variables that satisfy the regularity assumptions. Also suppose that $x_i \geq \Delta\ (\forall\ i = 1, \ldots, n)$. Then any incentive compatible, individually rational mechanism that maximizes expected revenues, subject to inequalities (3.2) and (3.3), and without allowing a reserve price, has purchase probabilities of $\mathbf{y}'^{\Delta}(\mathbf{v})$.

3.4 Analysis with a reserve policy and other variations

It is well understood that the use of a reserve price will generally increase a mechanism's expected revenues. In the present context, a reserve policy effectively means that the search provider does not carry out all value-enhancing trades. Instead, when constraint (3.1) is applied, a purchase probability of Δ is transferred from the lowest valuation firm to the highest valuation firm only when the difference between the firms' valuations is sufficiently large. Analogous to Myerson and Satterthwaite (1983, Theorem 4), "trade" occurs based upon a comparison of the highest virtual valuation (calculated as a "buyer" firm) and the lowest virtual cost (calculated as a "seller" firm). When constraint (3.2) is applied, a purchase probability of Δ is transferred from the k^{th} lowest to the k^{th} highest valuation firm $(k = 1, \ldots, \left\lfloor \frac{n}{2} \right\rfloor)$ only when the difference between the firms' valuations is sufficiently large.

Theorems 1 and 2 include, in their assumptions, that $x_i \geq \Delta\ (\forall\ i = 1, \ldots, n)$. This assumption means that the purchase probability of each firm is sufficiently large that all of the purchase probability Δ that is transferred to the highest valuation firm can come entirely at the expense of the purchase probability of the lowest valuation firm. If that assumption is not satisfied in an application of Theorems 1 and 2, the probabilities $\mathbf{y}^{\Delta}(\mathbf{v})$ would be modified so that the purchase probability Δ is taken from two or more of the low valuation firms. Similarly, if that assumption is not satisfied in an application of Theorems 1′ and 2′, the probabilities $\mathbf{y}'^{\Delta}(\mathbf{v})$ would be modified so that the purchase probability Δ is taken from more of the low valuation firms, which might also reduce the number of high valuation firms for which the purchase probability can be increased.

3.5 Discussion

Let us make a general observation about the information needed by the mechanism designer. The descriptions of $\mathbf{y}^{\Delta}(\mathbf{v})$ and of the basic mechanism that yields nonnegative expected revenues for $n \geq 3$ firms are entirely distribution free. So, in major respects, the results of Theorems 1 and 2 do not require knowledge of the distribution, $F(\bullet)$. However, adjustments to the mechanism to optimize fully the expected revenues are distribution dependent. For example, the determination of constant terms that optimize the expected revenues in Theorem 2—or that make the expected revenues nonnegative for $n = 2$ firms in Theorem 1—requires knowledge of the distribution. Similarly, determination of the optimal reserve policy depends on the distribution.

One additional, curious observation to make is that a reserve policy is usually associated with social loss, inasmuch as it prevents surplus-enhancing trades from occurring. However, in the present context, "trades" mean that interventions were made to responses so as to make one purchase probability lower than warranted by the AI model's innate information and to make another purchase probability higher than warranted. As such, a reserve policy is associated with less intervention to organic responses (i.e., providing more accurate responses) and so may be viewed as socially beneficial. One caveat to this observation is that, in a wider sense, the reserve policy could be reducing welfare, in that it reduces the shifting of purchases from firms that would value the purchase less to firms that would value the purchase more (i.e., it may be reducing producer surplus, which is one component of social surplus).

4. Inputs

In Section 2, artificial intelligence was used largely to generate responses to users' search requests. However, to realize the full potential of artificial intelligence in the context of sponsored search, we should also utilize an AI model to enable advertisers to express their bids more efficiently and effectively. Such a methodology may be utilized to generate traditional sponsored search results comprising ordered lists of sponsored links, as well as to generate nontraditional search results comprising unordered paragraphs of free-form prose or other novel outputs.

In traditional sponsored search auction systems, advertisers submit bids for keywords or concepts. Bids for keywords are well suited to search engines in which the user requests themselves comprise only a few search terms. However, bids for keywords are less tailored to search engines in which the user requests are more nuanced and written in conventional prose (and less likely to contain standard keywords).

For example, compare the three following search requests (all variations on the running example that we started in Section 2):

(1) Annapolis hotel.
(2) Mr. and Mrs. Jones, a conservative couple in their 50s, without children, are planning to spend 24 hours in Annapolis, Maryland. Can you write one page of text, recommending how Mr. and Mrs. Jones should spend 24 hours in Annapolis, Maryland? The text should include recommendations on things to do, places to eat, and places to stay.
(3) Mr. and Mrs. Jones, a conservative couple in their 50s, without children, are planning to spend 48 hours in Annapolis, Maryland. Can you write one page of text, recommending how Mr. and Mrs.

Jones should spend 48 hours in Annapolis, Maryland? The text should include recommendations on things to do, places to eat, and places to stay.

All three of these requests would be of interest to a hotel or motel located in Annapolis; and, likely, a hotel or motel located in Annapolis would like to bid for all three requests. However, there are three basic issues with traditional keyword auctions. First, depending on the sophistication of the search engine, it is unclear if request #2 or request #3 would necessarily trigger obvious keyword choices such as "Annapolis" AND "hotel" (since the requests do not contain the word "hotel").[10] Second, and more fundamentally, while the second and third requests are likely to trigger exactly the same keyword bids—their only difference is "24" in request #2 versus "48" in request #3—it is evident from their plain meaning that most Annapolis hotels would want to bid substantially higher for request #3 than for request #2. Third, above and beyond any information known about the user computer (such as its geographic location), the words contained in the requests could indicate substantially different values for different advertisers that may routinely bid for the same keywords. Consider the Historic Inns of Annapolis, which may be oriented toward couples and may not be oriented toward children. They may be willing to bid $3 for request #1, which is completely generic. However, the second and third requests indicate that the search is more likely to be associated with the hotel's target clientele, and so a click is more likely to convert to a booking. Furthermore, if they value a two nights' stay twice as much as they value a one night's stay, they may be willing to bid twice as much for the third request as for the second (e.g., $6 for request #2 and $12 for request #3).

In our ongoing example of Annapolis hotels, rows 2 – 4 of Table 3, below, illustrate bidding information that could be submitted by the Historic Inns of Annapolis (which we assume is oriented toward couples without children), while Rows 5 – 9 of Table 3 illustrate bidding information that could be submitted by the Rodeway Inn (which we assume is oriented toward families with children). The idea of submissions of bidding information is that they should provide bids for sufficiently many examples so that an intelligent bidding system would be able to infer values for most possible scenarios. For example, with the bidding information of Table 3, if the actual request had been "How should a couple without children spend 72 hours in Annapolis?", the AI model would infer that the Historic Inns of Annapolis and the Rodeway Inn would be willing to bid amounts of something like $18 and $3, respectively—since the actual request now encompasses a three-nights' stay.

[10] As noted in the Introduction, even at the time of the initial writing of this paper, a Google search on "kids shoes" and on "shoes for kids" (fully synonymous phrases) yielded different sponsored hyperlinks, arranged in different orders.

Advertiser	Request	Bid
Historic Inns of Annapolis	Annapolis AND hotel	$3
Historic Inns of Annapolis	How to spend 24 hours in Annapolis by a couple without children	$6
Historic Inns of Annapolis	How to spend 48 hours in Annapolis by a couple without children	$12
Rodeway Inn	Annapolis AND hotel	$2.50
Rodeway Inn	How to spend 24 hours in Annapolis by a family with children	$4
Rodeway Inn	How to spend 48 hours in Annapolis by a family with children	$8
Rodeway Inn	How to spend 24 hours in Annapolis by a couple without children	$1
Rodeway Inn	How to spend 48 hours in Annapolis by a couple without children	$2

TABLE 3: Set of bidding information from advertisers

To construct a system for processing bids in this way, we would begin with an artificial intelligence (AI) model that is pre-trained and fine-tuned with "organic" information (i.e., using the actual available data, without knowledge of any bids). Unlike the system used in Section 2, this system probably does not require the general knowledge base of ChatGPT. Instead, it needs to have high facility with language, it needs to have high facility with some basic reasoning, and it needs to be fast. Periodically (e.g., daily or hourly), a new instance of the AI model (from the point before being exposed to any bidding information) is loaded and it is fine-tuned with the entire set of current bidding information from advertisers. Once the AI model has been fine-tuned with the current bidding information, every time that a search request is received, the AI model is instructed to generate a list of the "leading advertisers" for the request (i.e., the advertisers who would be willing to bid the most for the request) and the amounts that they would be willing to bid for it.

The resulting list of leading advertisers and their synthetic bid amounts can be used as inputs, for whatever the auction process and the search outputs may be. If the outputs will be a traditional ordered list of clickable hyperlinks, then the auction system can run a Generalized Second Price auction on the synthetic bid amounts to determine the positions of sponsored links and their prices. If the outputs will be untraditional paragraphs of free-form text, then the auction system can run something closer to the approach described in Section 2.

Discussion

The specification of bids by expressing bidding information for a few insightful exemplary requests and then allowing artificial intelligence to extrapolate to all other requests appears to be a very powerful approach. It seems potentially to be a vast improvement over naming bids for specific keywords—and the

approach likely has wider application to other auction systems. However, the approach could also suffer from credibility issues, given that the company providing the AI system has a strong financial incentive to overestimate what advertisers would be willing to pay for requests. One approach for mitigating the credibility issue is as follows. An advertiser begins by submitting bidding information (such as rows 2 – 4 or rows 5 – 9 of Table 3). The AI system responds by providing the advertiser with a list of the most common requests that relate to the advertiser's submitted bidding information, together with the amounts that it estimates the advertiser would be willing to bid for these requests. The advertiser then has the choice of accepting the estimates in this list or of modifying estimated bid amounts using the advertiser's own numbers. This process can then be iterated until the advertiser and the AI system converge on a final list of acceptable bid amounts—and this final list is the one that the AI system applies as the advertiser's actual bidding information. Of course, in addition, the synthetic bid amounts applied by the AI system can themselves be made subject to verification and audit after the fact.

5. Conclusions

This paper has speculated about the likely form of sponsored search as results evolve from ordered lists of URLs to AI-based responses exemplified by ChatGPT. The speculation has been refined by applying some basic mechanism design, where the applied constraints are motivated by avoiding excessive intervention in users' purchase probabilities. When the constraint is that the sum of all deviations from the "organic" purchase probabilities must not exceed some fixed amount, the provider's expected revenues are optimized by increasing just one advertiser's purchase probability and by decreasing just one advertiser's purchase probability.

In Section 2, we also considered a mechanism in which the purchase probabilities of several advertisers were all increased in proportion to their bid amounts. Notwithstanding the mechanism design analysis, there is one scenario where multiple winners could still be optimal. The simple model assumes that all n advertisers are in competition for the same recommendation by the provider. In fact, many searches may yield recommendations for different products that are not in direct competition with one another: for example, the response to the "24 hours in Annapolis" search would highlight restaurants/bars, places of lodging, and sightseeing destinations. The optimal mechanism for such a scenario would probably not increase just one advertiser's purchase probability, but instead, would probably increase one advertiser's probability *in each purchase category*. Rather than trying to divide advertisers into categories—an exercise that is itself fairly arbitrary—and only increase the probability for one advertiser, it may be preferable to increase several advertisers' probabilities without regard to category. Determining whether this is the case will require additional research.

References

Cornes, R.C. and R. Hartley (2005), "Asymmetric Contests with General Technologies," *Economic Theory*, Vol. 26, pp. 923–946.

Cramton, P., R. Gibbons and P. Klemperer (1987), "Dissolving a Partnership Efficiently," *Econometrica*, Vol. 55, pp. 615–632.

Edelman, B., M. Ostrovsky, and M. Schwarz (2007) "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, Vol. 97, pp. 242–259.

Fallows, D. (2005), "Search Engine Users," Pew Internet & American Life Project of the Pew Research Center.

Federal Trade Commission (2017), "Blurred Lines: An Exploration of Consumers' Advertising Recognition in the Contexts of Search Engines and Native Advertising," FTC Staff Report, Dec. 2017.

Gallagher, D. (2025), "AI's Threat to Google Just Got Real," *The Wall Street Journal*, 10 May 2025.

Grant, N. and C. Metz (2022), "A New Chat Bot Is a 'Code Red' for Google's Search Business," *The New York Times*, Dec. 21, 2022.

Hansen, E. (2002), "FTC Wants Paid Search to Shape Up," *CNET*, Aug. 2, 2002.

Morris, C. (2025), "Elon Musk, longtime defender of open-source AI, is bringing advertising into his rogue Grok chatbot," *Fortune*, 8 August 2025.

Murgia, M., C. Criddle and G. Hammond (2024), "OpenAI explores advertising as it steps up revenue drive," *Financial Times*, 2 December 2024.

Murphy, H. (2025), "Elon Musk to introduce ads to X's AI chatbot," *Financial Times*, 7 August 2025.

Myerson, R.B. and M.A. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trade," *Journal of Economic Theory*, Vol. 29, pp. 265–281.

Palumbo, A. (2025a), "Alphabet Stock Sinks on Apple Testimony About Falling Search Queries," *Barron's*, May 7, 2025.

Palumbo, A. (2025b), "Advertising Could be AI's Next Frontier. Get Ready for Ads in Chatbots." *Barron's*, 18 June 2025.

Tibken, S. (2013), "FTC to Search Engines: Make it Clear Which Results are Ads," *CNET*, June 25, 2013.

Varian, H. (2007), "Position Auctions," *International Journal of Industrial Organization*, Vol. 25, pp. 1163–1178.

Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, L. Kaiser, and I. Polosukhin (2017), "Attention is All You Need," 31st Conference on Neural Information Processing Systems (NIPS 2017). URL https://arxiv.org/abs/1706.03762.

Vojnović, M. (2016), *Contest Theory: Incentive Mechanisms and Ranking Methods*. New York: Cambridge University Press.