AI & Data Obfuscation: Algorithmic Competition in Digital Ad Auctions

By Francesco Decarolis, Gabriele Rovigatti, Michele Rovigatti, Ksenia Shakhgildyan*

Two interrelated trends characterizing today's digital economy are the growing use of AI algorithms (AIAs) for pricing and other economic decisions, and the changes in data flow driven by the various initiatives for privacy protection enacted by both governments (for instance, through the GDPR in Europe and CCPA in the US) and the large digital platforms (like Apple's restrictions on tracking with the launch of IOS14 in April 2021 or the Google's decision to eliminate third party cookies).[1] Although the ongoing debate about regulating digital platforms is looking at both phenomena, it is doing so as if they were separate. But data are the key fuel of AIAs. Thus, any change to the type and quality of available data has an impact on the type and performance of the AIAs. This, in turn, implies that the large digital platforms might have incentives to strategically alter data flows to their advantage. Although the latter problem has been discussed by prominent literature in economics, its interactions with AIAs have not.[2]

This study intends to fill this gap by analyzing the case of digital advertising. Specifically, we want to understand whether a platform that sells ad slots on its search engine (as Google, Bing, or Yandex do) can benefit from obfuscating the data available to the AIAs bidding in its search auctions. Sponsored search is the most lucrative portion of digital advertising (worth about 40 percent of all digital ad revenues) and an area where algorithmic bidding, often via AIAs, is the norm. These algorithms require data to optimize bids, budgets, and keyword selection, but, to a large extent, it is the selling platform that determines the type, amount, frequency, and coarseness of the data released. As a result, platforms can potentially control the effectiveness of AIAs. Regulations like the forthcoming DMA in Europe will mandate that large platforms should disclose data, but it is far from easy to inform regulators on which type of data to focus on, especially given the conflicting role played by the privacy initiatives.[3] In this paper, we offer some clear evidence of the risk that platforms strategically obfuscate data through a series of simulated experiments where asymmetric bidders employ AIAs to compete in the Generalized second-price (GSP) auctions.[4] In this context, different amount of available information leads to differences in terms of possible states on which the bidders condition future bids. We find that when more detailed information is available to train the algorithms, the advertisers' rewards are higher, and conversely, the auc-

* Francesco Decarolis, Bocconi University, *francesco.decarolis@unibocconi.it*, Gabriele Rovigatti, Bank of Italy, *gabriele.rovigatti@bancaditalia.it*, Michele Rovigatti, Bocconi University, *mic.rovigatti@gmail.com*, Ksenia Shakhgildyan, Bocconi University, *ksenia.shakhgildyan@unibocconi.it*. The views expressed in the article are those of the authors and do not involve the responsibility of the Bank of Italy.

[1] Regarding the first trend, see (Chen, Mislove and Wilson 2016), (Competition and Authority 2018), (Calvano et al. 2020), (Assad et al. 2020), (Klein 2021), (Brown and MacKay 2021) and (Mehta and Perlroth 2023). Regarding the second trend, see (Alcobendas, Kobayashi and Shum 2021), (Aridor, Che and Salz 2022) and (Lefrere et al. 2022).

[2] See (Bergemann and Bonatti 2018) for an extensive introduction to markets for information, including an overview of the design and price of information for the case of sponsored search auctions.

[3] See, for instance, the conflicting provisions in the DMA that on the one hand enhance transparency toward advertisers (Art. 6g), while on the other place restrictions on targeted/micro-targeted ads (Art. 6aa).

[4] In terms of the methodology, our approach follows (Calvano et al. 2020) by setting up a series of computational experiments where AI algorithms set bids.

tioneer revenues decline. In particular, in our setting, the decision not to reveal competitor bids increases the platform revenues by 22%. Moreover, algorithmic bidding has a tendency to sustain low bids relative to the competitive benchmark case.

I. Generalized Second-price Auction

Consider three advertisers $i \in \{1, 2, 3\}$ bidding in an online advertising auction with two slots. Advertisers value a click on their ad differently: $v_1 = 3$, $v_2 = 2$, $v_3 = 1$. If an ad is placed on the first slot, it gets five clicks, whereas the second slot leads to two clicks. Denote these click-through-rates (CTRs): $x^1 = 5$, $x^2 = 2$. We discretize the set of feasible bids $\mathcal{B}$ on the interval $[B_{\min,} B_{\max}] = [0.2, 3]$, with $k = 15$ possible bids so that the step between the bids is 0.2. The bid of advertiser i is denoted b_i. The first slot is assigned to the highest bidder and the second to the second highest. When several bids are equal, the slots are allocated randomly. Denote the rank of advertiser i's bid $\rho(i)$, then, the resulting payoff is $v_i x^{\rho(i)}$. In the GSP mechanism, each bidder pays the price-per-click equal to the bid of the advertiser placed below him. As a result, bidder i's reward in the GSP auction can be written as $\left(v_i - b^{\rho(i)+1}\right) x^{\rho(i)}$.

The static GSP auction has many Nash equilibria. For this reason, (Varian 2007) and (Edelman, Ostrovsky and Schwarz 2007) introduced a refinement of the set of equilibria, the *lowest-revenue locally envy-free equilibrium* (EOS), which is predominantly used in the literature on the GSP as a competitive benchmark.[5] The EOS equilibrium of the three-player game is given by $b_1 > b_2$, $b_2 = 1.6$, $b_3 = 1$, and leads to auctioneer revenue $R = 10$.

These studies model the GSP auction as a full information game because, back in the days when they were written, the amount of feedback that was available to bidders was extensive. Bidders had auction-level information such that it was as if they knew other players' bids. Below we discuss how the data that search engines pass to advertisers has evolved over time becoming coarser. This tendency is what we refer to as data obfuscation.

II. Design Features of the AIAs Experiments

The auction game described above is repeated many times. This repetition is what allows the AIAs to learn through a process of trial and error how to optimize bids in order to maximize a reward. In particular, the AIAs that we consider are Q-learning algorithms bidding against each other and learning simultaneously.[6] Their training entails striking a balance between exploration (trying out new strategies) and exploitation (using the obtained knowledge).

The knowledge

The knowledge of each algorithm is represented by the *Q-matrix*, which is the matrix of expected rewards from each possible bid in each possible state of the game. For each bidder i, in each period t, it is $Q_t^i(s, b)$, where $b \in \mathcal{B}$, and $s \in \mathcal{S}$. Here, the states of the game can contain different amounts of information about the past auction outcomes. For example, in one of the experiments, each state is defined by the previous bids of all players. In the first period, each cell of the Q-matrix is initialized randomly.

The experimentation

To fully explore the Q-matrix, the algorithm should visit different actions in different states, even the ones that it finds not

[5] A Nash equilibrium is locally envy-free if $x^{\rho(i)}(v_i - b^{\rho(i)+1}) \geq x^{\rho(i)-1}(v_i - b^{\rho(i)})$ for every i. EOS refinement is the lowest-revenue Nash equilibrium which satisfies this condition. This refinement is especially important because it conforms with the search engines' tutorials on how to bid in these auctions. See, for instance, the Google AdWord tutorial in which Hal Varian teaches how to maximize profits by following this bidding strategy: `https://www.youtube.com/watch?v=tW3BRMld1c8`. As EOS showed, such equilibria induce the same allocations and payments as truthful bidding in the VCG, and they are fully characterized by the following conditions: denote by S the number of available slots, then $b_1 > b_2$, $b_i = v_i$ for all $i > S$, and for all $i = 2, \ldots, S$, $b_i = v_i - \frac{x^i}{x^{i-1}} (v_i - b_{i+1})$.

[6] This is the approach pioneered by (Calvano et al. 2020). For an overview of this type of AIAs see (Sutton and Andrew 2018).

optimal, given prior knowledge. We use an ϵ-*greedy* exploration strategy. At each iteration, the algorithm chooses the bid that currently leads to the highest value of Q-matrix in a given state with a probability $1-\epsilon$, and with probability ϵ chooses a random bid among all possible ones.[7] We will use a declining with time exploration rate $\epsilon_t = e^{-\beta * t}$, where $\beta > 0$ is the annihilation coefficient of the exploration rate.[8]

Among the many features that characterize how an AIA is designed, two play a particularly key role in our analysis. These are the updating rule and the data usage.

II.A. The Updating Rule

The information obtained in period t is used for updating the Q-matrix. The algorithm starts from an initial Q-matrix. After choosing bid b_t^i in state s_t, the algorithm observes r_t^i as well as s_{t+1}, and updates $Q_t^i(s,b)$ for $s = s_t$ and $b = b_t^i$. The updating can happen in a number of ways. We consider two main approaches that the literature describes as synchronous and asynchronous updating rules:

I. *Asynchronous Updating*: For each i, $Q^i(s_t, b_t^i)$ is updated using the following "temporal difference" update rule:

$$Q_{t+1}^i(s_t, b_t^i) = (1-\alpha)Q_t^i(s_t, b_t^i) + \alpha * (r_t^i(s_t, b_t^i, b_t^{-i}) + \delta * \max_{b' \in \mathcal{B}} Q_t^i(s', b')).$$

Here α is the learning rate.[9] The learning rate determines to what extent the new information substitutes the old ("how much" the algorithm learns from new bids and received rewards). At the same time, $Q_{t+1}^i(s,b) = Q_t^i(s,b)$ for all $s \neq s_t$, $b \neq b_t^i$. Asynchronous updating only requires knowledge of the reward received from the submitted bid.

II. *Synchronous Updating*: $Q^i(s_t, b)$ is updated for all bids $b \in \mathcal{B}$ with a reward $r_t^i(s_t, b, b_t^{-i})$ that the bidder would have received had it submitted a different bid, given the bids of other players:

$$Q_{t+1}^i(s_t, b) = (1-\alpha)Q_t^i(s_t, b) + \alpha * (r_t^i(s_t, b, b_t^{-i}) + \delta * \max_{b' \in \mathcal{B}} Q_t^i(s', b')),$$

and $Q_{t+1}^i(s,b) = Q_t^i(s,b)$ for all $s \neq s_t$. Thus, synchronous updating requires calculating the rewards for all the bids of i, $r_t^i(s_t, b, b_t^{-i})$, hence also for those bids that were not submitted $b \neq b_t^i$.

From the description above it is clear that the feasibility of the two approaches above crucially hinges on the available data and on how such data can be used to calculate the counterfactual reward associated with actions that are not taken. In our setting, the GSP auction has very clear rules to determine the allocation of slots and payments depending on the bids received.[10] Hence, if bidder i observes b^{-i}, the calculation of his reward under any possible b^i holding fixed b^{-i} is trivial and, hence, using a synchronous updating rule is feasible. But absent data on b^{-i}, this counterfactual calculation is unfeasible. In this case, an asynchronous updating rule is instead possible because its implementation requires observing only the reward associated with the bid effectively submitted by i.

[7] In what follows we assume that if several bids lead to the same value of Q-matrix in a given state, one of them is chosen randomly. Results hold also when the algorithms are conservative and choose the lowest bid, and they are presented in Appendix A.

[8] We use $\beta = 8.137e-07$, for the simulations with 15,000,000 iterations, and $\beta = 1.22e-05$ for the simulations with 1,000,000 iterations, so that at the last iteration, the probability of exploration is only 0.00005.In particular, $exp(-15,000,000 * 8.137e-07) = 0.00005$ and $exp(-1,000,000 * 1.22e-05) = 0.00005$.

[9] In what follows, $\alpha = 0.1$ which is a standard in computer science literature. We have also explored other learning rates, and our results are qualitatively similar. Moreover, $\delta = 0.95$ unless stated otherwise.

[10] A synchronous algorithm thus operates in a way similar to how economists approach equilibrium analysis. See (Asker, Fershtman and Pakes 2022) for a discussion.

II.B. The Data Usage

The data is used not only to calculate the rewards associated with the different actions but also to keep track of the state. In this regard, there are two polar cases that we can consider in terms of how the data is used:

I. *Stateful algorithms*: Stateful algorithms maintain a record of previous bids and use this information to inform their decisions.

II. *Stateless algorithms*: Stateless algorithms, on the other hand, do not retain information from previous steps. They make decisions based solely on the current reward.

A key difference between stateful and stateless algorithms is that the former, since it has memory, can respond differently to the same received rewards, based on the previous state. This, in turn, allows for dynamic strategies that are not possible with stateless algorithms. But data requirements are greater for stateful algorithms if they need to keep track of past bids, relative to the case of stateless algorithms which do not require such information.

III. Data Policy: Obfuscation Stategies by the Platform

A search engine hosting GSP auctions has ample latitude about the data that it releases to bidders. Different considerations, from the technological feasibility to economic features like the reputation that the platform has (or seeks to establish) will inform its data policy. For the sake of clarity, we will focus on two extreme scenarios (*Full Information* vs *No information*), but also, among the infinitely many cases in between, present an intermediate case (*Partial Information*). These three cases differ in terms of what the platform reveals about bids:

I. *Full Information*: In every period, the bidder observes not only the current reward but also the bids of the other players submitted in the past period.

II. *Partial Information*: In every period, the bidder observes not only the current reward but also her bid submitted in the past period and the price paid.

III. *No information*: The only information that the bidder observes is the reward she received after submitting a particular bid.

Before relating these three cases to examples of data policies adopted by platforms, let us connect them to the feasibility of different AIAs designs. It is only under full information that Stateful synchronous algorithms are feasible. Instead, under no information, Stateless asynchronous is the only possible form of AIAs. Partial information is incompatible with the use of synchronous AIAs but allows asynchronous AIAs to retain some, limited form of memory containing one's own bids (we will refer to this type of AIAs as Partial asynchronous algorithms). Hence, the platform data policy determines to a significant extent the type of AIAs that advertisers can use and a movement away from full information is what we consider a data obfuscation strategy.

We can now turn to a few examples of data policies and their evolution. We focus on the case of Google. As mentioned early, (Varian 2007) and (Edelman, Ostrovsky and Schwarz 2007) who pioneered the equilibrium analysis of the search auctions decided that a complete information game was an adequate approximation of the environment faced by the bidder. This choice was in sharp contrast with the canonical auction literature but was motivated by the specificities of the environment. As stated in (Varian 2007): *"(...) one might ask how likely it is that advertisers know what they need to know to implement a full information equilibrium. (...) Google reports click and impression data on an hour-by-hour basis and a few days of experimentation can yield pretty good estimates of the number of clicks received for different bids. Furthermore, Google itself offers a "Traffic Estimator" that provides an esti-*

mate of the number of clicks per day and the cost per day associated with the advertiser's choice of keywords. Finally, third-party companies known as "Search Engine Managers (SEMs)" offer a variety of services related to managing bids. The availability of such tools and services, along with the ease of experimentation, suggest that the full-information assumption is a reasonable first approximation. As we will see below, the Nash equilibrium model seems to fit the observed choices well." Fast forwarding to 2021, the situation is radically changed to the point that many SEMs have been thrown out of business by how Google has limited their access to the data and, among those still in business, several have abandoned those activities that involved bidding on clients' behalf.[11] The "Traffic Estimator" tool has also been dismissed.

Among the changes to the data policy of Google, one that has received substantial attention from the industry is that involving the "search terms report". This report used to be a crucial tool to assess how each single keyword ad was performing. But starting in September 2020, Google modified it to contain exclusively "terms that a significant number of users searched for, even if a term received a click." Industry specialists have argued that this led to at least 20 percent of search terms becoming invisible for advertisers.[12] Another related instance occurred in July 2021 when Google announced "changes to phrase match and broad match modifiers": following changes that had begun in February 2021, the new system in place from July 2021 meant that keyword matches became broader, thus making it harder for advertisers to relate the keyword (for which they bid) to the user queries (which, if generating clicks, trigger the advertiser payments).[13]

Other examples exist and all describe the same pattern toward data obfuscation.[14] Hence, despite the richness of the potentially available information, sponsored search auctions now release so little information to advertisers that even the most extreme scenario that we labeled above as the *No information* one is likely a reasonable approximation of how this market currently works. Indeed, as much as it might appear absurd that an advertiser does not even know the price it paid for bidding on a keyword, this is what the combination of a second price system (that decouples bids from prices) and the revised data policies described above produce: a keyword bid gets applied through broad matches to multiple queries in ways that advertisers cannot control anymore neither ex ante (due to the broad match modifiers) nor ex-post (due to the revised search terms report). This might also explain why Google is moving toward a system of data-driven attribution.[15]

Different forces might be behind the revised data policies described above. For instance, regarding the changes to the search terms report, Google motivated it with the aim "to maintain our standards of privacy and strengthen our protections around user data." The same privacy narrative is behind all of the recent revisions to data policies adopted by the large platforms.[16] With-

[11]For a discussion of the role of the SEMs in the period of their active engagement in bidding and the potential threat it represented for the platform revenues see (Decarolis, Goldmanis and Penta 2020), (Decarolis and Rovigatti 2021), and (Decarolis et al. 2021).

[12]See panel (A1a) and (A1b) in Figure A1 in the Online Appendix reporting screenshots of the Google announcement and of a news article on its effects.

[13]See panel (A1c) in Figure A1 in the Online Appendix reporting a screenshot of the Google announcement.

[14]For instance, this is the case of the elimination of the average position. To know in which slot an ad was shown, the average position used to be a highly informative metric. However, a data policy change in February 2019 replaced it with coarser information describing what percent of ads appear at the top of the page (and at the very top of the page). See panel (A1d) in Figure A1 in the Online Appendix reporting a screenshot of the Google announcement.

[15]Attribution is a fundamental element in digital advertising: it relates a conversion to an action (in our case, a bid on a keyword). In September 2021, Google announced changes to the default attribution of clicks to bids offering a cryptical description of it: "data-driven attribution uses the Shapley value solution concept from cooperative game theory". See details in Figure A2 in the Online Appendix.

[16]Apple, for instance, described privacy protection as the reason for its IOS14 "do not track" feature. But, privacy implications aside, this revision turned out to

out questioning the underlying intentions, it is nevertheless important to assess what the consequences, intended or not, might be. This is what we do in the next section by evaluating how AIAs perform under different information regimes.

IV. Results

We begin by contrasting the outcomes under the two most extreme scenarios: *Full Information* vs *No information.* In terms of the AIAs, this means contrasting the algorithm that exploits the most data, the *Stateful Synchronous* algorithm which requires *Full Information* to the algorithm that requires the least data, the *Stateless Asynchronous* algorithm, the only type of algorithm implementable under *No information.*

Table 1 presents these baseline results. The first row summarizes the outcomes when the bidders have access to the information on the competitors' bids and use *Stateful Synchronous* algorithms, whereas the second row shows the outcomes when the platform restricts access to the information on the competitors' bids, and as a result, the advertisers use *Statelss Asynchronous* algorithms. For each of the settings, we ran the experiment 50 times. Column 1 reports the average bids at convergence for each of the players in the decreasing order of valuations.[17] Each run differs only in terms of the randomly initialized Q-matrix, as well as random exploration. Column 2 presents the average rewards at convergence, while column 3 reports the average auctioneer revenue across runs as well as its 95% confidence interval for the revenue (in squared brackets).

The decision not to reveal competitor bids increases the platform's average revenues by 22% from 7.2 to 8.76. That is driven by the reduction in the reward of the highest-value player from 9 to 7.46 due to the increase in the bid of the second-highest-value player, from 1.2 to 1.51.

To expand our analysis, we consider a series of other experimental designs, starting from the one that applies to the scenario of *Partial Information.* Table 2 illustrates the results for this case in its first row where we consider an AIA that is asynchronous but has a memory of the past price paid by the agent. Consistent with the intuition, this *Partial Asynchronous* algorithm produces average auctioneer revenue lower than in case of *No Information* (when *Statelss Asynchronous* algorithms are used), but still significantly higher than the average auctioneer revenues under the *Full Information* (when *Stateful Synchronous* algorithms are used instead). In particular, average auctioneer revenue increases by 8% from 7.2 to 7.79 when the auctioneer instead of revealing all the bids, only reveals the bid of the competitor who is placed right below a given advertiser, and as a result defines a price paid by this advertiser.

Two other cases are considered and reported in the last two rows of Table 2. These two cases look at how the performance of the *Stateful Synchronous* algorithm changes if we either shut down the forward-looking element of the payoff (i.e., we set δ=0 in the updating rule) or if we eliminate its memory of past bids (i.e., adopt a stateless algorithm). These two cases are of limited practical relevance within our setting as, if the platform discloses the full information needed by the synchronous algorithm, there would be no reason to ignore past states or future rewards. But they do serve an important illustrative role to explain our findings. Indeed, in both of these alternative versions of the synchronous algorithm, the bids at convergence are very close to those of the baseline *Stateless Asynchronous* and so are the revenues. These latter results are suggestive that the behavior observed for the baseline *Stateful Synchronous* algorithm crucially depends on the possibility of adopting dynamic strategies, which indeed require both memory of the past and attention to the future.

greatly benefit the Apple ad network at the expense of the Facebook one, see `https://www.ft.com/content/074b881f-a931-4986-888e-2ac53e286b9d`.

[17]The advertiser with value-per-click $v_1 = 3$ on average at convergence bids 2.03, the advertiser with value-per-click $v_2 = 2$ bids 1.2, whereas the advertiser with $v_3 = 1$ bids 0.6 when *Stateful Synchronous* algorithms are played.

Before exploring why dynamic strategies are important, let us complete the description of the baseline findings. So far we have focused on the bids and revenues at convergence; that is, on what the algorithms do once they have attained stable behavior.[18]But convergence requires a large number of periods.[19] Figure 1 shows the evolution of the auctioneer's revenues, individual bids, and rewards (vertical axis) by the percent or the total number of iterations (horizontal axis). The black line represents the mean of the distribution of revenues across simulation runs, the dark grey zone is the area between the 25th and 75th percentiles, while the light grey is the zone between the 10th and 90th percentiles. The algorithms start to coordinate long before convergence is achieved. The auctioneer revenues start from a fairly large value, but this is simply because the algorithms initially randomize uniformly across bids that, on average, lead to a revenue similar to the EOS level. This effect disappears as experimentation starts to be less prominent, and eventually, auctioneer revenues converge to the lower level. As can be seen from Panels (d)-(i), most of the difference in revenues between *Stateful Synchronous* and *Stateless Asynchronous* experiments is driven by the reduction in the reward of the highest-value player due to the increase in the bid of the second-highest-value player.

IV.A. Drivers of the Baseline Findings

Although it is notoriously difficult to fully open up the black box of how AIAs work, in our case a key driver of the behavior of the *Stateful Synchronous* algorithm relative to the *Stateless Asynchronous* algorithm can be traced back to dynamic strategies. Building upon our earlier discussion of the experiments in Table 2, here we compare the bid evolution when the synchronous algorithm learns to bid with different discounting of the future payoffs. In particular, Figure 2 shows the evolution of the smoothed bids (moving average was applied) by the percent or the total number of iterations in the baseline *Stateful Synchronous* experiment with $\delta = 0.95$ as well as in the one with $\delta = 0$. The difference is striking. We can observe that while with $\delta = 0.95$ when one of the players increases the bid, the other right away follows, no such behavior is observed in the case when $\delta = 0$. Moreover, as discussed earlier, the average auctioneer revenues in the case of Stateful Synchronous experiments with $\delta = 0$ are 8.57 and not statistically different from the ones in Stateless Asynchronous experiments.

That logic can be seen clearly in Figure 3. Here, we have focused on the *Stateful Synchronous* algorithm's bids at convergence. We start from the final Q-matrix, but then introduce an exogenous shock to the bid of Player $v_2 = 2$. Instead of bidding 1.2, she deviates to bidding 1.6 in the period that we call 0. What follows is that when Player $v_2 = 2$ deviates, Player $v_3 = 1$ increases her bid (punishes with a higher price).[20] Importantly, just after a few periods, all players return to equilibrium bidding. [21]

V. Generalizations

In this section, we extend our results along three dimensions. First, we consider alternative settings of AIAs hyperparameters; second, we look at variations in the GSP auction game; lastly, we consider alternative auction designs. In general, all of the extensions below lead to the same qualitative outcome of the baseline findings, with differences only in the magnitude of the revenue increase via obfuscation.

Starting from the case of alternative AIAs, we consider different designs of the Q-learners. In one experiment, we consider alternative bid selection methods for the cases in which different actions are associated with identical Q-values. In partic-

[18]See Appendix C for the definition and discussion of convergence.

[19]Much less for the Stateless algorithms, but on the order of millions for the Stateful.

[20]That is the result of a long history of learning, not a random action chosen by Player $v_3 = 1$ on a wider interval of actions since Q-value of the bid 1.2 could have been updated just by an insignificant amount.

[21]The exact number of periods depends on a particular simulation run, and varies from 2 to 5.

ular, the main results also hold when the algorithms are conservative and choose the lowest bid, and they are presented in Appendix A. In another experiment, we consider asymmetric grids, spanning the same bid space, but featuring a different number of actions for each player.

The second extension regards variations to the GSP stage game. In Appendix D, we show that the increase in auctioneer revenues due to the information restriction can be much higher than in our baseline experiment. To do that, we run an experiment with three asymmetric advertisers bidding in an online advertising auction with three slots, taken from (Milgrom and Mollner 2014). Advertisers value a click on their ad differently: $v_1 = 15$, $v_2 = 10$, $v_3 = 5$, so the ratio of the values is the same as in our baseline setting. What is different is the relative click-through rates. If an ad is placed on the first slot, it gets 100 clicks, whereas the second slot leads to three clicks, and the third to one click. In this case, we find that the increase in auctioneer revenues due to the data obfuscation is 78%.

The third and last variation involves alternative auction designs. In Appendix E, we consider another mechanism, namely the Vickrey-Clarke-Groves (VCG), widely used in online advertising.[22] We find that for the VCG, the decision not to reveal competitor bids increases the platform's average revenues by 38%. Moreover, the auctioneer revenues under the VCG setting tend to be lower than those under the GSP. The results are presented in Table A5. This latter finding is complementary to the results of a growing number of studies on the performance of AIAs across auction formats.[23]

VI. Conclusions

In this study, we analyzed the role of AI in digital ad auctions. Through computational experiments, we evaluated the performance of bidding algorithms powered by AI across several stylized auction games characterized by different information levels. We find that when more detailed information is available to the algorithms, the advertisers' rewards tend to be higher, and conversely, the auctioneer revenues tend to decline. In particular, the decision not to reveal competitor bids increases the platform revenues by 22%. That increase is driven by the fact that with *Limited Information* (the bidder observes only her own reward) only *Statelss Asynchronous* algorithms can be used, while in case of *Full Information* (the bidder observes not only the current reward but also the bids of the other players submitted in the past period), bidders could instead exploit *Stateful Synchronous* algorithms. Moreover, algorithmic bidding sustains low bids under the GSP relative to the competitive benchmark. The results are robust to a number of extensions, notably the change of the auction format from GSP to VCG.

These results highlight several important features regarding how AI might shape the working of ad auctions. Moreover, our findings might help to understand recent decisions by the dominant selling platforms. In the context of search advertising, advertisers on Google have experienced a reduction in the amount of type of data available to optimize their bidding strategies. This deliberate increase in the extent of *data fogginess*, while responding to the growing privacy concerns among search users, creates a trade-off with market competition.

Generalize back: all platforms.

New competition issue: data is an input to vertical partners, foreclosure incentive. This shall guide data regulations, with an eye on AI

In a follow-up study, we also analyze the tension created by the deployment of AI bidding tools by the platform. These tools

[22] VCG is allegedly used by Facebook.

[23] For instance, a thorough analysis of first price vs second price auctions in the presence of stateless algorithms is conducted in (Banchio and Skrzypacz 2022). This study considers a setting with symmetric players. The authors find no difference between the auctioneer revenues for the case of a symmetric second-price auction with 2 players. For the case of the first-price auction, instead, stateless synchronous algorithms are less likely to converge on collusive outcomes. This is also coherent with (Asker, Fershtman and Pakes 2021) who find that in the context of Bertrand competition, a stateless synchronous algorithm restores competition.

might be trained by the platform on data that is superior to that available to advertisers or their intermediaries, thus leading such platform-sponsored tools to outperform the competition. But once all the bidding activity is directly delegated to the platform itself, the risks for competition or, at the very least, the lock-in effect on advertisers are significant.

REFERENCES

Alcobendas, Miguel, Shunto Kobayashi, and Matthew Shum. 2021. "The Impact of Privacy Measures on Online Advertising Markets." *Working Paper.*

Aridor, Guy, Yeon-Koo Che, and Tobias Salz. 2022. "The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR." *Rand Journal of Economics*, forthcoming.

Asker, John, Chaim Fershtman, and Ariel Pakes. 2021. "Artificial Intelligence and Pricing: The Impact of Algorithm Design."

Asker, John, Chaim Fershtman, and Ariel Pakes. 2022. "The Impact of AI Design on Pricing." *Journal of Economics and Management Strategies*, forthcoming.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2020. "Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Working paper.*

Banchio, Martino, and Andrzej Skrzypacz. 2022. "Artificial Intelligence and Auction Design." *Available at SSRN 4033000.*

Bergemann, Dirk, and Alessandro Bonatti. 2018. "Markets for Information: An Introduction." *Coweles Foundation Discussion Paper N. 2142.*

Brown, Zach, and Alexander MacKay. 2021. "Competition in Pricing Algorithms." *American Economic Journals: Microeconomics*, forthcoming.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2020. "Artificial Intelligence, Algorithmic Pricing, and Collusion." *American Economic Review*, 110(10): 3267–97.

Chen, L., A. Mislove, and C. Wilson. 2016. "An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace." *Proceedings of the 25th International Conference on World Wide Web*, 1339–1349.

Competition, and Markets Authority. 2018. "Pricing algorithms: Economic working paper on the use of algorithms to facilitate collusion and personalised pricing." *CMA Working Paper CMA94.*

Decarolis, Francesco, and Gabriele Rovigatti. 2021. "From Mad Men to Maths Men: Concentration and Buyer Power in Online Advertising." *American Economic Review*, forthcoming.

Decarolis, Francesco, Maris Goldmanis, and Antonio Penta. 2020. "Marketing Agencies and Collusive Bidding in Online Ad Auctions." *Management Science*, Vol.66, No 10, 4433-4454.

Decarolis, Francesco, Maris Goldmanis, Antonio Penta, and Ksenia Shakhgildyan. 2021. "Bid coordination in sponsored search auctions: Detection methodology and empirical analysis."

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz. 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review*, 97(1): 242–259.

Klein, Timo. 2021. "Autonomous algorithmic collusion: Q-learning under sequential pricing." *The RAND Journal of Economics*, 52(3): 538–558.

Lefrere, Vincent, Logan Warberg, Cristobal Cheyre, Veronica Marotta, and Alessandro Acquisti. 2022. "Does Privacy Regulation Harm Content Providers? A Longitudinal Analysis of the Impact of the GDPR." *Working paper.*

Mehta, Aranyak, and Andres Perlroth. 2023. "Auctions without commitment in the auto-bidding world." *arXiv preprint arXiv:2301.07312.*

Milgrom, Paul, and Joshua Mollner. 2014. "Equilibrium selection in multiplayer games with auction applications." Citeseer.

Sutton, Richard S., and G. Barto Andrew. 2018. *Reinforcement Learning: An Introduction.* MIT Press.

Varian, Hal. 2007. "Position auctions." *International Journal of Industrial Organization*, 25(6): 1163–1178.

Table 1—: Limit Bids, Rewards, and Auctioneer Revenues: Baseline Experimental Designs

GSP Auction with Three Asymmetric Players			
	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2
			[7.03, 7.37]
Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76
			[8.39, 9.13]

Note: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

Table 2—: Limit Bids, Rewards, and Auctioneer Revenues: Other Experimental Designs

Three Asymmetric Players (GSP)			
	Bids	Individual Rewards	Revenue
Partial Asynchronous	$(2.22, 1.34, 0.58)$	$(7.83, 2.79, 0.12)$	7.79
			[7.28, 8.29]
Stateful Synchronous ($\delta = 0$)	$(2.46, 1.47, 0.61)$	$(7.64, 2.79, 0.0)$	8.57
			[8.06, 9.08]
Stateless Synchronous	$(2.49, 1.49, 0.6)$	$(7.55, 2.8, 0.0)$	8.65
			[8.31, 8.99]

Note: UPDATE THIS NOTE XXXX: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

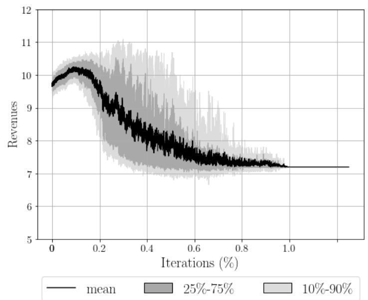

(a) Stateful Sync Revenues

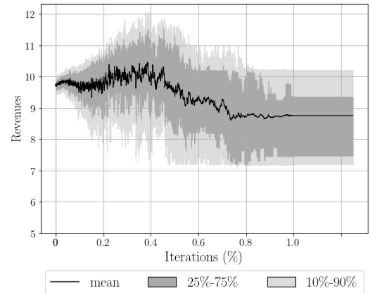

(b) Stateless Async Revenues

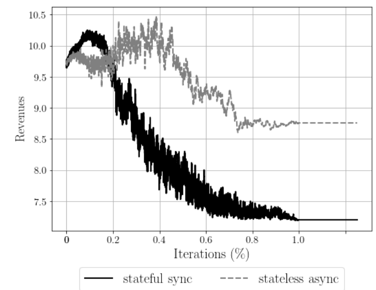

(c) Revenues Evolutions

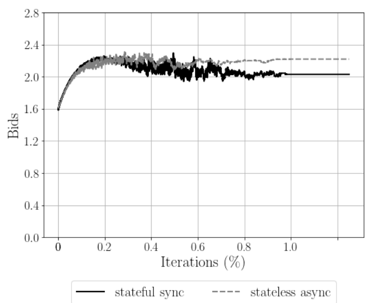

(d) Bids of Player $v_1 = 3$

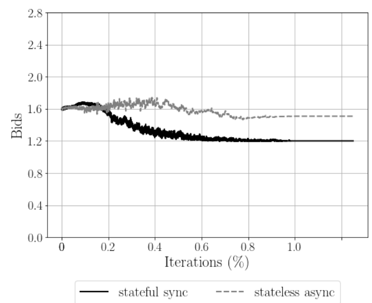

(e) Bids of Player $v_2 = 2$

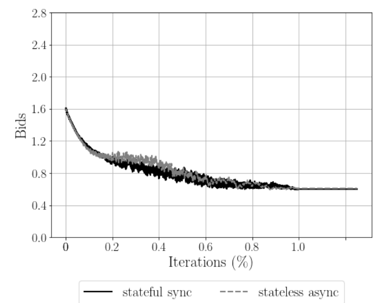

(f) Bids of Player $v_3 = 1$

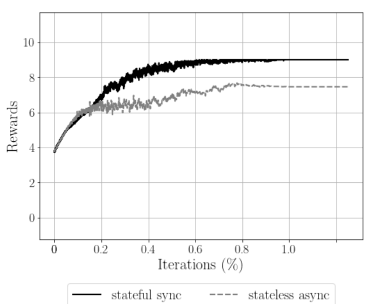

(g) Rewards of Player $v_1 = 3$

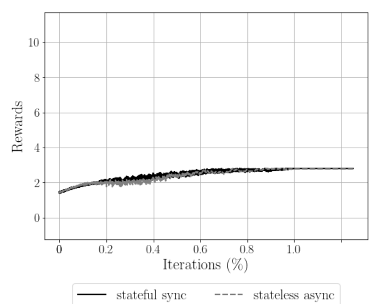

(h) Rewards of Player $v_2 = 2$

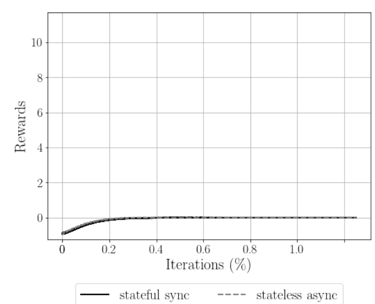

(i) Rewards of Player $v_3 = 1$

Figure 1. : Evolution of Auctioneer's Revenues, Bids, and Rewards

Note: Panels (a)-(c) show the evolution of the auctioneer's revenues (vertical axis) by the percent or the total number of iterations (horizontal axis). The black line represents the mean of the distribution of revenues, the dark grey zone is the area between the 25th and 75th percentiles, while the light grey is the zone between the 10th and 90th percentiles. Panels (d)-(i) show the evolution of the individual bids and rewards (vertical axis) by the percent or the total number of iterations (horizontal axis). The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

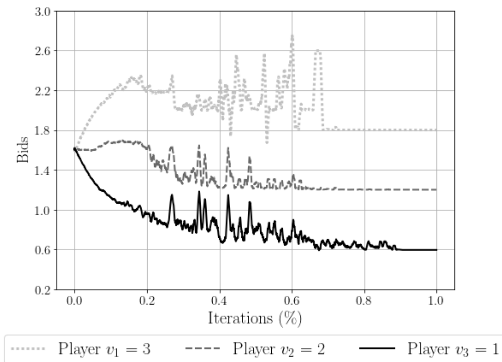

(a) Bids when $\delta = 0.95$

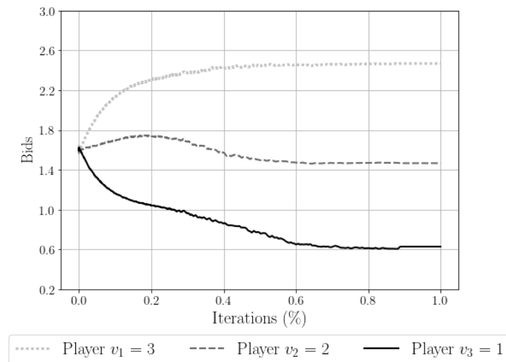

(b) Bids when $\delta = 0$

Figure 2. : Evolution of Bids in a Single Run

Note: Panel (a) shows the evolution of the smoothed bids (moving average was applied) by the percent or the total number of iterations in just one run of the baseline Stateful Synchronous experiment. Panel (b), instead shows the evolution of the smoothed bids (moving average was applied) by the percent or the total number of iterations in just one run of Stateful Synchronous experiment with $\delta = 0$. Stateful Sync experiment was run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, and both $\delta = 0.95$, and $\delta = 0$.

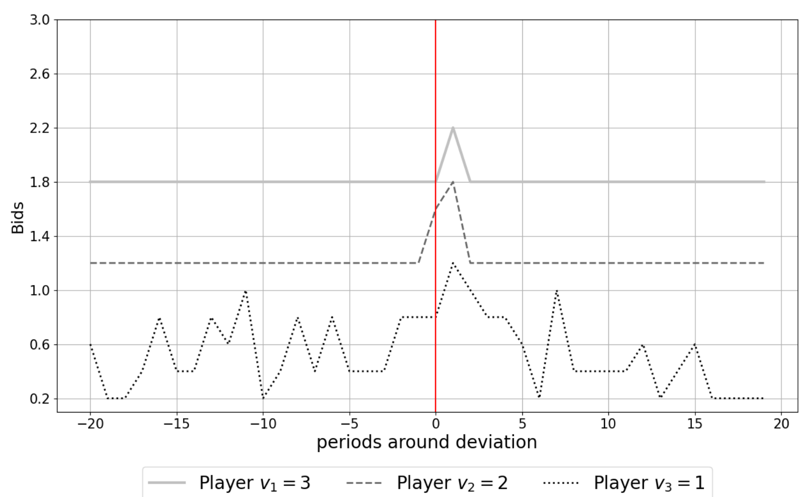

Figure 3. : Evolution of Bids in a Single Run of when Player $v_2 = 2$ deviates at period 0

Note: The figure shows the evolution of the bids by the iteration from the moment of the forced deviation of Player $v_2 = 2$ to raise his bid to 1.6 instead of his bid 1.2 at convergence of the Stateful Synchronous algorithm.

Appendix: For Online Publication Only

A. Conservative algorithms

In this appendix, we present results similar to the ones in Table 1 and Figure 1 with the only difference that in the exploration process, the algorithms are conservative and choose the smallest bid among the ones that lead to the highest value of the Q-matrix in a given state. See Table A1 and Figure A3.

The decision not to reveal competitor bids increases the platform's average revenues by 23% from 6.4 to 7.86. That is driven by the reduction in the reward of the highest-value player from 9 to 7.49 due to the increase in the bid of the second-highest-value player, from 1.2 to 1.49.

B. Comparison with literature

For comparison with the literature, we also consider the *Stateless Synchronous* and *Stateful Asynchronous* algorithms. Results are presented in Table A2 below. If we were to focus on the comparison between Stateless algorithms, one with Asynchronous Updating, and the other with Synchronous, our findings extend the ones by (Banchio and Skrzypacz 2022). We find no statistically significant difference between the auctioneer revenues for the case of asymmetric GSP.[1] Importantly, in order to use Synchronous Updating, the information on the bids of other players should be available. And given the availability of such information, players would naturally retain at least a short-term memory of it, and condition on it. As soon as we implement the *Stateful Synchronous* algorithms instead of the *Stateless*, the results are reversed, and coordination is again sustained by the algorithms. In Table A2 we also present results for the *Stateful Asynchronous* experiment that is the one used by (Calvano et al. 2020). We do not find a statistically significant difference with the *Stateless* algorithms.

C. Convergence

Since the algorithmic environments are not stationary, there is no guarantee that the algorithms converge. Even if the algorithms converge, the basic question is whether the bids converge toward the Nash Equilibrium (NE) predictions. Still, convergence can be verified ex-post. Following (Calvano et al. 2020), we use the following definition: convergence is deemed to be achieved if for each player the optimal strategy does not change for 100,000 consecutive periods for the case of *Stateful Synchronous* experiments. That is, if for each player i and each state s the action $b_{i,t}(s) = argmax[Q_{i,t}(b, s)]$ stays constant for 100,000 repetitions, we assume that the algorithms have completed the learning process and attained stable behavior. We stop the session when this occurs, and in any case after 15 million iterations. In the *Stateless Asynchronous* experiment, 1,000,000 iterations were enough for the convergence to be reached with enough exploration since the Q-matrix, in this case, takes the form of Q-vector since the state is a singleton, thus the convergence check was based on 1,000 iterations. Only a very few runs didn't converge, and thus for all the charts and tables, we considered only converged runs as in (Calvano et al. 2020) and (Banchio and Skrzypacz 2022).

For the *Stateful Synchronous*, 15 million iterations are required for the probability of exploration to decrease to 0.00005, so that convergence can be reached: $exp(-15,000,000 * 8.137e-07) = 0.00005$. If the rival is experimenting at even a 1% rate, the environment is still too non-stationary for the algorithm to converge. As a result, convergence is achieved only when experimentation is nearly terminated. In turn, such a small β leads to Q-matrix exploration that is sufficient for the exper-

[1] (Banchio and Skrzypacz 2022) consider the second-price auctions with symmetric players, as well as the first-price. The authors find no difference between the auctioneer revenues for the case of a symmetric second-price auction with 2 players. For the case of the first-price auction, instead, stateless synchronous algorithms are less likely to converge on collusive outcomes. Also, (Asker, Fershtman and Pakes 2021) find that in the context of Bertrand competition, a stateless synchronous algorithm restores competition.

iments with 3 players. As can be seen in Table A3, each cell of the Q-matrix is visited at least 72 and 428 times. In some simulations, the algorithms experience cyclical behavior and do not bid a constant amount.

D. Milgrom and Mollner (2014) example

Consider the case of three asymmetric advertisers $i \in \{1, 2, 3\}$ bidding in an online advertising auction with three slots taken from (Milgrom and Mollner 2014). Their valuations are $v_1 = 15$, $v_2 = 10$, $v_3 = 5$, respectively, while the click-through rates amount to $x^1 = 100$, $x^2 = 3$ and $x^3 = 1$. We discretize the set of feasible bids $\mathcal{B}$ on the interval $[B_{\text{min},} B_{\text{max}}] = [1, 15]$, with $k = 15$ possible bids so that the step between the bids is 1. The EOS equilibrium in this case is given by $b_1 > b_2$, $b_2 = 9.8$, $b_3 = 3.3$, and leads to auctioneer revenue $R = 990$.

The results of the experiment, reported in table A4, show that the magnitude of the increase in auctioneer revenues due to the information obfuscation strictly depends on the structure of the auction prizes. In this case, it amounts to a 78% increase.

E. Comparison with the VCG

In table A5, we compare the GSP (columns 1 to 3) and the VCG (columns 4 to 6) mechanisms. We find that for both auction designs the decision not to reveal the competitor bids increases the platform's average revenues. Moreover, the auctioneer revenues under the VCG setting tend to be lower than those under the GSP setting.

Table A1—: Limit Bids, Rewards, and Auctioneer Revenues with conservative AIs

GSP Auction with Three Asymmetric Players			
	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(2.07, 1.2, 0.2)$	$(9.0, 3.6, 0.0)$	6.4 [6.4, 6.4]
Stateless Asynchronous	$(2.14, 1.49, 0.2)$	$(7.49, 3.59, 0.0)$	7.86 [7.56, 8.16]

Note: This table presents the same results as in Table 1 with the only difference that in the exploration process, the algorithms are conservative and choose the smallest bid among the ones that lead to the highest value of the Q-matrix in a given state. Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

Table A2—: Limit Bids, Rewards, and Auctioneer Revenues under various Experimental Designs

Three Asymmetric Players (GSP)			
	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2 [7.03, 7.37]
Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76 [8.39, 9.13]
Stateless Synchronous	$(2.49, 1.49, 0.6)$	$(7.55, 2.8, 0.0)$	8.65 [8.31, 8.99]
Stateful Asynchronous	$(2.17, 1.43, 0.8)$	$(6.83, 2.32, 0.16)$	8.92 [8.34, 9.5]

Note: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

Table A3—: Summary Statistics on Count Matrices

	count	mean	std	min	25%	50%	75%	max
State Visits	3375.0	64701.54	543513.35	1741.5	2322.56	4726.5	11719.12	15970379.25
Cell Visits	50625.0	4313.44	36229.21	116.1	154.78	315.1	781.75	1064691.95

Note: The table reports summary statistics on the average number of visits counted in each Q-matrix state and in each Q-matrix cell. for the Stateful Synchronous experiments. Specifically, count matrices are averaged across runs for every player. Then, they are reshaped as a vector of dimension $[1 \times k]$, where k is the size of the q-matrix. This vector is player specific, and it is used for computing summary statistics at the player level. For the analysis of states, average count matrices are first summed row-wise so as to get a vector of dimension $[S \times 1]$, where S is the number of states of the q-matrix, which is then used to compute summary statistics on states at the player-level.

Table A4—: Limit Bids, Rewards, and Auctioneer Revenues under various Experimental Designs

Three Asymmetric Players (GSP)			
	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(11.45, 4.3, 2.17)$	$(1069.35, 23.57, 5.0)$	437.02 [416.05, 458.0]
Stateless Asynchronous	$(13.36, 7.72, 2.01)$	$(728.0, 23.96, 5.0)$	778.04 [711.29, 844.8]

Note: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

Table A5—: Comparison of the GSP and VCG

Three Asymmetric Players						
	GSP			VCG		
	Bids	Individual Rewards	Revenue	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2 [7.03, 7.37]	$(2.52, 1.21, 0.6)$	$(10.18, 2.8, 0.0)$	6.02 [5.63, 6.41]
Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76 [8.39, 9.13]	$(2.79, 1.94, 0.62)$	$(7.94, 2.76, 0.0)$	8.3 [7.81, 8.79]

Note: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

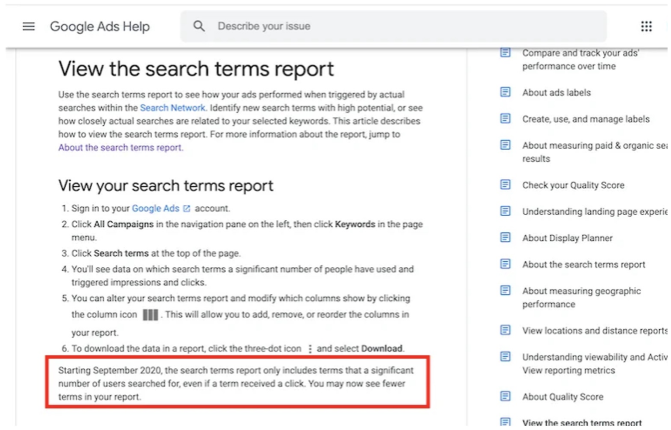

Google Ads Help

Describe your issue

View the search terms report

Use the search terms report to see how your ads performed when triggered by actual searches within the Search Network. Identify new search terms with high potential, or see how closely actual searches are related to your selected keywords. This article describes how to view the search terms report. For more information about the report, jump to About the search terms report.

View your search terms report

1. Sign in to your Google Ads account.
2. Click **All Campaigns** in the navigation pane on the left, then click **Keywords** in the page menu.
3. Click **Search terms** at the top of the page.
4. You'll see data on which search terms a significant number of people have used and triggered impressions and clicks.
5. You can alter your search terms report and modify which columns show by clicking the column icon . This will allow you to add, remove, or reorder the columns in your report.
6. To download the data in a report, click the three-dot icon and select **Download**.

Starting September 2020, the search terms report only includes terms that a significant number of users searched for, even if a term received a click. You may now see fewer terms in your report.

- Compare and track your ads' performance over time
- About ads labels
- Create, use, and manage labels
- About measuring paid & organic sea results
- Check your Quality Score
- Understanding landing page experie
- About Display Planner
- About the search terms report
- About measuring geographic performance
- View locations and distance reports
- Understanding viewability and Activ View reporting metrics
- About Quality Score
- View the search terms report

(a) Search Term Report

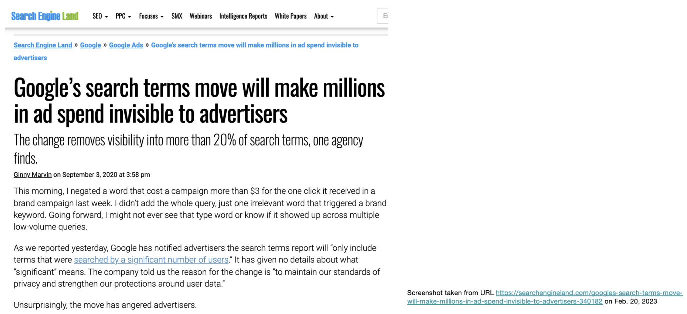

Search Engine Land SEO PPC Focuses SMX Webinars Intelligence Reports White Papers About

Search Engine Land » Google » Google Ads » Google's search terms move will make millions in ad spend invisible to advertisers

Google's search terms move will make millions in ad spend invisible to advertisers

The change removes visibility into more than 20% of search terms, one agency finds.

Ginny Marvin on September 3, 2020 at 3:58 pm

This morning, I negated a word that cost a campaign more than $3 for the one click it received in a brand campaign last week. I didn't add the whole query, just one irrelevant word that triggered a brand keyword. Going forward, I might not ever see that type word or know if it showed up across multiple low-volume queries.

As we reported yesterday, Google has notified advertisers the search terms report will "only include terms that were searched by a significant number of users." It has given no details about what "significant" means. The company told us the reason for the change is "to maintain our standards of privacy and strengthen our protections around user data."

Unsurprisingly, the move has angered advertisers.

Screenshot taken from URL https://searchengineland.com/googles-search-terms-move-will-make-millions-in-ad-spend-invisible-to-advertisers-340182 on Feb. 20, 2023

(b) Impacts of the Search Term Report Change

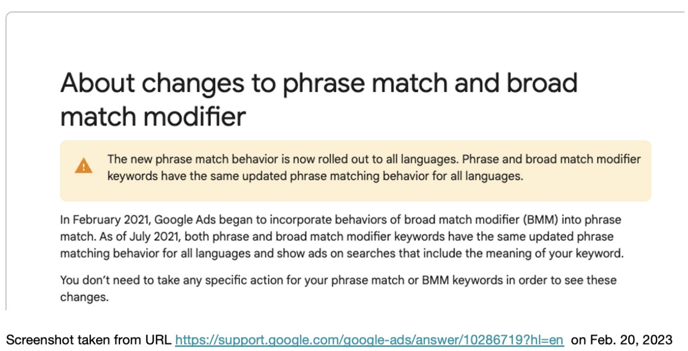

About changes to phrase match and broad match modifier

The new phrase match behavior is now rolled out to all languages. Phrase and broad match modifier keywords have the same updated phrase matching behavior for all languages.

In February 2021, Google Ads began to incorporate behaviors of broad match modifier (BMM) into phrase match. As of July 2021, both phrase and broad match modifier keywords have the same updated phrase matching behavior for all languages and show ads on searches that include the meaning of your keyword.

You don't need to take any specific action for your phrase match or BMM keywords in order to see these changes.

Screenshot taken from URL https://support.google.com/google-ads/answer/10286719?hl=en on Feb. 20, 2023

(c) Broad Match Modifiers

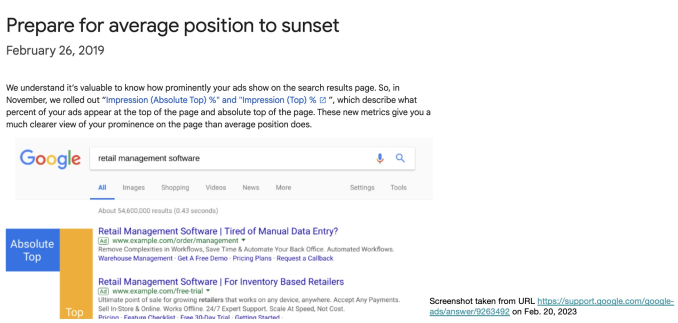

Prepare for average position to sunset

February 26, 2019

We understand it's valuable to know how prominently your ads show on the search results page. So, in November, we rolled out "Impression (Absolute Top) %" and "Impression (Top) % ", which describe what percent of your ads appear at the top of the page and absolute top of the page. These new metrics give you a much clearer view of your prominence on the page than average position does.

Screenshot taken from URL https://support.google.com/google-ads/answer/9263492 on Feb. 20, 2023

(d) Position Information

Figure A1. : Examples of Data Policy Changes

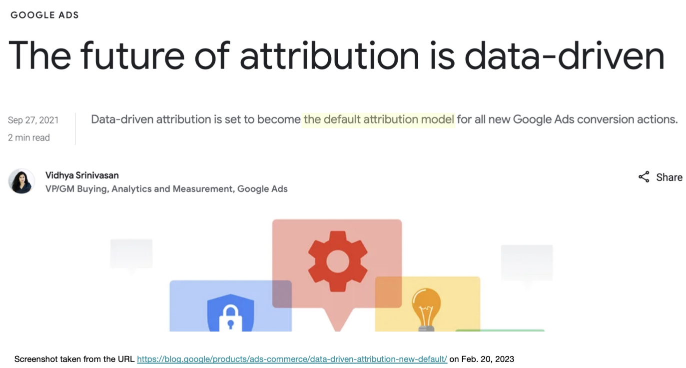

(a) Attribution

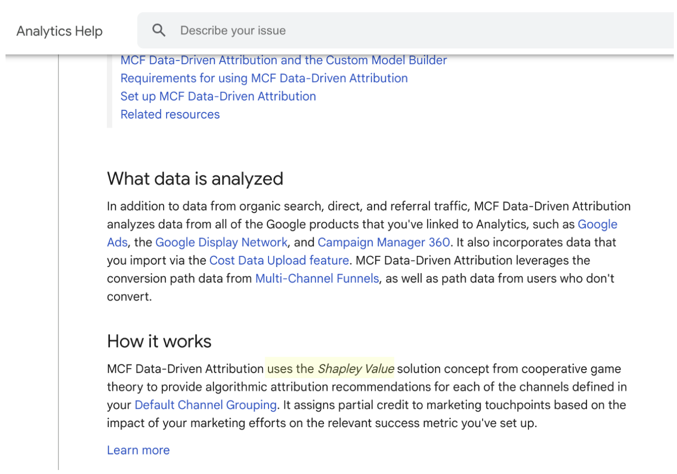

Analytics Help

Describe your issue

MCF Data-Driven Attribution and the Custom Model Builder
Requirements for using MCF Data-Driven Attribution
Set up MCF Data-Driven Attribution
Related resources

What data is analyzed

In addition to data from organic search, direct, and referral traffic, MCF Data-Driven Attribution analyzes data from all of the Google products that you've linked to Analytics, such as Google Ads, the Google Display Network, and Campaign Manager 360. It also incorporates data that you import via the Cost Data Upload feature. MCF Data-Driven Attribution leverages the conversion path data from Multi-Channel Funnels, as well as path data from users who don't convert.

How it works

MCF Data-Driven Attribution uses the *Shapley Value* solution concept from cooperative game theory to provide algorithmic attribution recommendations for each of the channels defined in your Default Channel Grouping. It assigns partial credit to marketing touchpoints based on the impact of your marketing efforts on the relevant success metric you've set up.

Learn more

(b) Details on the Attribution Model

Figure A2. : Changes in Default Attribution

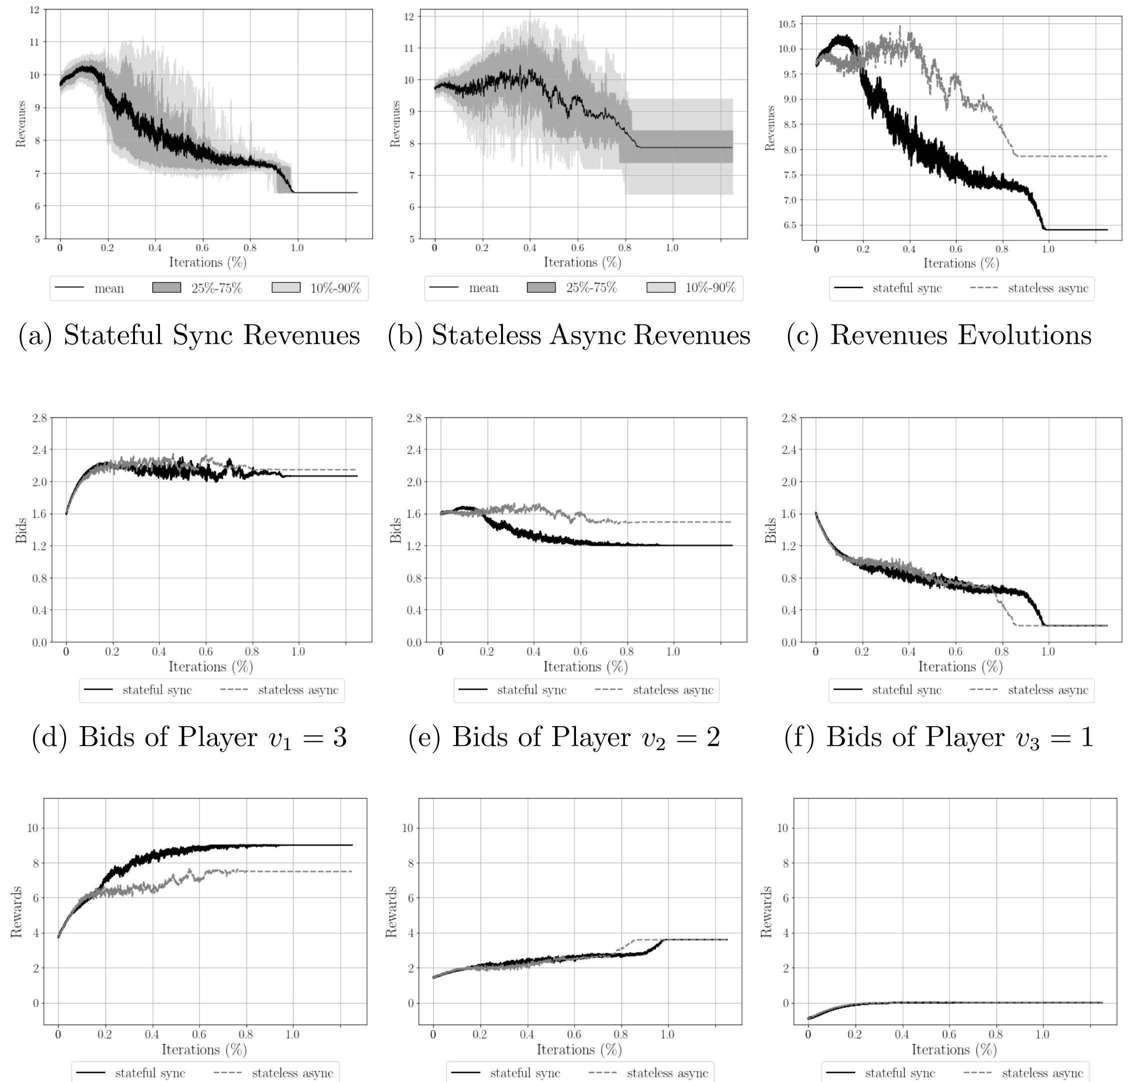

(a) Stateful Sync Revenues (b) Stateless Async Revenues (c) Revenues Evolutions

(d) Bids of Player $v_1 = 3$ (e) Bids of Player $v_2 = 2$ (f) Bids of Player $v_3 = 1$

(g) Rewards of Player $v_1 = 3$ (h) Rewards of Player $v_2 = 2$ (i) Rewards of Player $v_3 = 1$

Figure A3. : Evolution of Auctioneer's Revenues, Bids, and Rewards

Note: Panels (a)-(c) show the evolution of the auctioneer's revenues (vertical axis) by the percent or the total number of iterations (horizontal axis). The black line represents the mean of the distribution of revenues, the dark grey zone is the area between the 25th and 75th percentiles, while the light grey is the zone between the 10th and 90th percentiles. Panels (d)-(i) show the evolution of the individual bids and rewards (vertical axis) by the percent or the total number of iterations (horizontal axis). The Stateful Sync experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Async experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.