Artificial Intelligence & Data Obfuscation: Algorithmic Competition in Digital Ad Auctions

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

Data aren't just the fuel of artificial intelligence. Data granularity, frequency, and quality determine the feasibility and performance of AI algorithms. In the context of the generalized second-price auction used to sell internet search ads, we conduct simulated experiments with asymmetric bidders competing through Q-learning algorithms under different information structures on rival bids. We find that when less detailed information is available to train algorithms auctioneer revenues are substantially and persistently higher. This underscores the incentive for the digital platforms designing data-sharing policies to distort data flows to their advantage by strategically obfuscating data.

Keywords: Asymmetric Information, Auctions, Procurement, Artificial Intelligence, Collusion, Data, Privacy, Data Governance, Digital Advertising, Competition, Digital Platforms.

JEL Codes: C73, D82, D83, D18, D44.

* 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. We gratefully acknowledge financial support from the ERC Consolidator Grant - CoDiM -#101002867.

I. Introduction

Two interrelated trends characterizing today's digital economy are the growing use of Artificial Intelligence Algorithms (AIAs) for pricing, bidding, and other economic decisions, and the restrictions on the data flows enacted by both governments and large digital platforms that are often officially motivated by privacy-related concerns. Although these two trends have been treated separately in both academic and policy debates, they are inextricably linked due to data being the key input of AIAs. But data is more than a fuel, it is a determinant of AIAs in the sense that any change to the type and quality of available data has an impact on the type and performance of the feasible AIAs. In this study, we explore the connection between data and AIAs and illustrate why, absent any consideration of privacy, a platform might have an incentive to reduce the data disclosed to alter the effectiveness of its business users' AIAs.

A growing literature in economics focuses on data, data markets, and data governance (see (Fainmesser, Galeotti and Momot 2023), as well as the reviews in (Bergemann and Morris 2019), (Bergemann and Bonatti 2019), and (Goldfarb and Que 2023)). However, the link between platform data strategies and the AIAs used by businesses operating through these platforms is still unexplored. This study contributes to filling this gap by analyzing the effects of different data policies on AIAs in the context of a market, digital advertising, that has all the main ingredients required. First, a few large digital platforms decide on the data that are fed to the advertisers (and their intermediaries) bidding in the online ad auctions where ad space is sold. Second, these data have recently become more coarse: as discussed in the next section, over the last few years, the main platform of this sector (Google) has implemented multiple revisions to its data policy that resulted in less and less information being provided to the bidders. Third, most of the bidding happens via algorithms, oftentimes powered by AI (see (IAB 2018)).

Specifically, we focus on sponsored search which is the most lucrative portion of digital advertising.[1] We want to understand whether a platform that sells ad slots on its search engine via Generalized second-price (GSP) auctions (as Google, Bing, or Yandex do) can impact the AIAs bidding in its auctions through the decision on what to reveal about keyword bids. These algorithms require data to optimize bids, budgets, and keyword selection, but 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 Digital Markets Act (DMA) mandate that ad platforms disclose data to advertisers, but there are no exact specifications on the type of data.

In terms of the methodology, our approach follows (Calvano et al. 2020) by setting up a series of computational experiments where AI algorithms interact in a repeated game. In our simulated experiments, asymmetric bidders (i.e., with different valuations) employ AIAs to compete in a GSP auction for a particular keyword ad. Different classes of AIAs are feasible, depending on the amount of available information: indeed, the available data determines the possible states of the world on which the players condition future bids as well as the possibility to evaluate the counterfactual scenarios. Hence, data policies determine the types of AIAs that can be used. We focus on the case of Q-learning algorithms to take advantage of the transparency of the learning process entailed by these algorithms.[2]

We show that, when less detailed information is available to train the algorithms, this restricts two key design features of the algorithm: the learning rule and the memory. Specifically, we focus on the information on the competitors' bids. This is motivated by a striking difference between the Yahoo! auction design, which used to reveal all the bids to the advertisers, and the current Google design, which keeps the bids private.[3] We show that when we move between a scenario of full

[1]See IAB, Internet Advertising Report, April 2021. According to Statista, ad spending in the US in the Search Advertising segment is projected to reach $133.50bn in 2023, well above the projected spending in the Banner Advertising segment ($66.81bn) or the Social Media Advertising segment ($94.42bn).

[2]Q-learning is the basis of all reinforcement learning algorithms, and the problem of bidding in the ad auctions is a reinforcement learning problem by nature. See, for example, "AI in Advertising: Real-Time Bidding & Reinforcement Learning."Moreover, Q-learning algorithms are often analyzed in the literature since they are characterized by a few parameters, each having a straightforward economic interpretation.

[3]See the discussion in (Lahaie 2006).

information where the platform releases data on all the bids submitted in the past period to one where no data is released, the advertisers' rewards decline and conversely, the auctioneer revenues increase. In particular, in our baseline setting, the decision not to reveal competitor bids increases the platform revenues by 22%.

The second main finding is that algorithmic bidding tends to sustain low bids relative to the competitive benchmark case. Furthermore, this tendency is stronger the more data are available.

We then explore the mechanisms behind our findings which are linked to dynamic strategies and the possibility of calculating counterfactuals. Through a series of controlled alterations of the baseline algorithms, we show how data on competitors' bids lead to more interaction, especially between the bidders with lower valuations. Interestingly, that does not lead the advertisers to bid more agressively, instead advertisers that get position bid lower to secure the lower price paid. Importantly, this insight would have not been attainable in a stylized case with symmetric bidders, as generally done in the literature, and this underscores the relevance of looking at more realistic environments reflecting the asymmetry that characterizes search auctions.

Finally, we also explore the generalizability of our experimental results. We consider three dimensions along which to extend our results. First, given the wide margins by which the design of AIAs can vary, we consider a series of alternative settings of Q-learning AIAs; second, we look at variations in the GSP auction game by analyzing alternative examples in the literature; lastly, we consider alternative auction designs, notably the Vickrey-Clarke-Groves (VCG) mechanism. In general, we find that all of the extensions lead to the same qualitative outcome of the baseline findings, with differences only in the magnitude of the revenue changes across different data obfuscation scenarios.

Literature Our study contributes to three main branches of the literature. While the relevance of information is an old idea, its modern incarnation in the form of studies on privacy (Acquisti, Taylor and Wagman 2016) and (Goldfarb and Que 2023), information design (Bergemann and Morris 2019), and information markets (Bergemann and Bonatti 2019) reflects the key role that data has come to play in today's digital economies. Online advertising, due to both its economic relevance and its fundamental connection to data markets, has been one the main areas where this literature has looked for applications, see (Bergemann et al. 2022*b*) and (Bergemann et al. 2022*a*). Although not looking at AI and focused on bidding in display ad auctions, the closest paper to ours is (Alcobendas, Kobayashi and Shum 2021) on how the elimination of third-party cookies, damages advertisers.

The second branch is that on the effects of AI and its impacts on pricing. (Agrawal, Gans and Goldfarb 2022) offers a general overview, while (Calvano et al. 2020) is the first study to explore how and why prices set by AIAs might differ from equilibrium predictions. How AIAs design features drive pricing has been stressed by (Asker, Fershtman and Pakes 2023). Most of the studies in this area adopt the method of simulated experiments that we also follow, but an important exception is (Assad et al. 2023) who use data on gas pricing in Germany to show how the deployment of AIAs in this market led to substantial price increases. This literature is rapidly growing in many directions,[4] including one that specifically focuses on AI bidding auctions, see (Dütting et al. 2017), (Heidekrüger et al. 2021), (Bichler, Fichtl and Heidekrüger 2021) and (Banchio and Skrzypacz 2022).

The third branch is that on online advertising and, more specifically, sponsored search. Starting with the seminal works of (Edelman, Ostrovsky and Schwarz 2007) and (Varian 2007), this literature has been extended along many directions by (Borgers et al. 2013), (Athey and Nekipelov 2014), (Gomes and Sweeney 2014), (Blake, Nosko and Tadelis 2015), (Che, Choi and Kim 2017), (Simonov, Nosko and Rao 2018), (Decarolis, Goldmanis and Penta 2020), (Simonov and Hill 2021), (Decarolis and Rovigatti 2021), (Deng et al. 2023). We contribute to this literature by introducing the analysis of AIAs. Search advertising also sits at the core of the antitrust concerns about digital platforms (Morton and Dinielli 2020). In this respect, our results drive attention to a series of changes in the data flows charac-

[4]See (Chen, Mislove and Wilson 2016), (Competition and Authority 2018), (Klein 2021), (Brown and MacKay 2023), (Banchio and Mantegazza 2022), (Lambin 2023) and (Mehta and Perlroth 2023).

terizing this sector and offer a rationale for why they might concern regulators.

II. Data Policy: Obfuscation Strategies by the Platform

A search engine hosting online search 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. We will focus on two extreme scenarios:

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. *Limited Information*: The only information that the bidder observes is the reward she received after submitting a particular bid.

These cases differ in terms of what the platform reveals about bids and are motivated by a striking difference between the auction designs observed in the market which range from the case of Yahoo! revealing all bids (Lahaie 2006), to that of Google currently keeping bids private. Interestingly, recent revisions to the Google data policy do not even allow any more advertisers to link bids to the clicks they triggered. We discuss in detail in Online Appendix A some of these cases, but here we offer a few examples of data policies and their evolution. Focusing on Google, (Varian 2007) and (Edelman, Ostrovsky and Schwarz 2007) argued that a complete information game was an adequate approximation of the environment faced by bidders. Indeed, the availability of such tools as a "Traffic Estimator" and services provided by third-party companies known as "Search Engine Managers (SEMs)," along with the ease of experimentation, suggested that the full-information assumption is a reasonable first approximation.

Fast forward to today, the situation has radically changed as Google now limits data access in several ways. One change receiving substantial attention from the industry is that involving the "search terms report". This used to be a crucial tool to assess how each keyword ad was performing. But starting in September 2020, it was modified 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 to advertisers. Subsequently, in July 2021 Google announced "changes to phrase match and broad match modifiers." In this case, 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).[5] Combined, these revisions happening in 2020 and 2021 mean that a keyword bid is assigned 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).

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 digital platforms.[6] Without questioning the motivations, it is important to assess the implications of how the ad auctions work. In our context, we will refer to the tendency of search engines to disclose increasingly coarser data to advertisers as data obfuscation.

III. 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

[5]See references and details in Online Appendix A and Figure A1.

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

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 $r^i = \left(v_i - b^{\rho(i)+1}\right) x^{\rho(i)}$.

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 of this full information game, the *lowest-revenue locally envy-free equilibrium* (EOS), which is predominantly used in the literature on the GSP as a competitive benchmark.[7] 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 $R = 10$.

IV. Design Features of the AIAs Experiments

In this section, we describe the design features of the AIAs and relate them to different information regimes chosen by an online advertising platform. 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 the expected present value of the reward stream.[8] In particular, the AIAs that we consider are Q-learning algorithms bidding against each other and learning simultaneously.[9] Their training entails striking a balance between exploration (trying out new strategies) and exploitation (using the obtained knowledge).

[7]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 Vickrey-Clarke-Groves (VCG) auction, 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 $i = 2, \ldots, S$, $b_i = v_i - \frac{x^i}{x^{i-1}}\left(v_i - b_{i+1}\right)$.

[8]Which equals for player i, $E[\sum_{t=0}^{t=\infty} \delta^t r_t^i]$, where δ is the discount factor.

[9]For an overview of this type of AIA see (Sutton and Andrew 2018).

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 the *Full Information* experiment setting, 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 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.[10] We use a declining with time exploration rate $\epsilon_t = e^{-\beta * t}$, where $\beta > 0$ is the annihilation coefficient.[11]

Among the many features that characterize how an AIA is designed, two play a particularly key role in our analysis: the updating rule and memory of the past.

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 the reward r_t^i as well as s_{t+1}, and updates $Q_t^i(s, b)$. The updating can happen in a number of ways. We consider two main approaches that the literature describes as synchronous and asynchronous updating rules:

[10]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, or greedy and choose the highest. See Appendix C.

[11]We use $\beta = 8.137e - 07$, for the simulations with 15 million iterations, and $\beta = 1.22e - 05$ for the simulations with 1 million iterations, so that at the last iteration, the probability of exploration is 5e-06.

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

$$Q^i_{t+1}(s_t, b^i_t) = (1-\alpha)Q^i_t(s_t, b^i_t) + \alpha * (r^i_t(s_t, b^i_t, b^{-i}_t) + \delta * \max_{b' \in \mathcal{B}} Q^i_t(s_{t+1}, b')).$$

Here α is the learning rate.[12] 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, $Q^i_{t+1}(s,b) = Q^i_t(s,b)$ for all $s \neq s_t$, $b \neq b^i_t$. 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 $r^i_t(s_t, b, b^{-i}_t)$ that the bidder would have received had it submitted a bid b, given the bids of other players b^{-i}:

$$Q^i_{t+1}(s_t, b) = (1-\alpha)Q^i_t(s_t, b) + \alpha * (r^i_t(s_t, b, b^{-i}_t) + \delta * \max_{b' \in \mathcal{B}} Q^i_t(s_{t+1}, b')),$$

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

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.[13] Hence, if bidder i observes the bids of competitors b^{-i}, the calculation of her reward under any possible b^i holding fixed b^{-i} is transparent and, therefore, using a synchronous updating rule is feasible. An asynchronous updating rule can be used instead even without data

[12] In what follows, $\alpha = 0.1$ which is a standard in computer science literature. We have also explored other learning rates, in particular, $\alpha = 0.05, 0.2, 0.3$, and our results are qualitatively similar. Moreover, $\delta = 0.95$ unless stated otherwise.

[13] See (Asker, Fershtman and Pakes 2023) for a discussion.

on b^{-i},since its implementation requires observing only the reward associated with the bid effectively submitted by i.

Memory

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 can be considered:

I. *Stateful algorithms*: Stateful algorithms maintain a record of previous bids and use this information to inform their decisions. In particular, in case of *Full Information*, we define for each of the players $s_t = (b^i_{t-1}, b^{-i}_{t-1})$.

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. In particular, in that case, $s_t = \emptyset$.

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. However, data requirements are greater for stateful algorithms as they need to keep track of past bids, relative to the case of stateless algorithms which do not require such information.

Link between Data Policies and AIAs

Finally, let us connect different data policies to the feasibility of different AIA designs. It is only under *Full Information* that *Stateful Synchronous* algorithms (Stateful algorithms with Synchronous updating rule) are feasible. Instead, under *Limited Information* the *Stateful Synchronous* can no longer be used, but the *Stateless Asynchronous* (Stateless algorithms with Asynchronous updating rule) can. The key point here is that 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. After presenting our baseline results, we consider extensions in terms of different information structures and the kinds of algorithms that can be used in those cases.

V. Results

Table 1 presents 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 *Stateless 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.[14] 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.

Before diving into the discussion of the main drivers of the results, let us describe the limit strategies of each of the players in individual runs of the experiments. In *Stateful Synchronous* experiments, Player $v_2 = 2$ (*middle-value* player) always converges to bidding 1.2. Player $v_1 = 3$ (*high-value* player) converges to a stable bid, which varies slightly across runs, but predominantly equals 1.8. Instead Player $v_3 = 1$ (*low-value* player), cycles over bids in the interval from the lowest possible bid 0.2 to the bid 1 (which is right below the bid of the second player at convergence), all of which lead to zero reward for the lowest player. In turn, in individual runs of the *Stateless Asynchronous* experiment, *middle-value* player

[14]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.

always converges to a stable bid, with some variation over the runs across the bids 1.4, 1.6 and 1.8, but predominantly either 1.6, or 1.4. *High-value* player also always converges to a stable bid, in most of the runs it is 2. As in *Stateful Synchronous* case, *low-value* player iterates between all bids from 0.2 to 1.

Next, we consider a series of alternative experimental designs. The lower panel of Table 1 illustrates the results for this case in its first row where we consider an AIA that is asynchronous but conditions on the past price paid by the agent. Consistent with the intuition, this *Partial Asynchronous* algorithm produces average auctioneer revenue lower than in the case of *Limited Information* (when *Stateless Asynchronous* algorithms are used), but still significantly higher than the average auctioneer revenues under the *Full Information* (when *Stateful Synchronous* algorithms are used instead): the revenue increases by 8%, from 7.2 to 7.79.

Two other cases are considered and reported in the last two rows of Table 1. These 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 because, 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 in explaining 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 both the memory of the past and attention to the future.

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.[15] But convergence requires a large number of periods.[16] Figure 1 shows the evolution of the auctioneer's

[15]See Appendix B for the definition and discussion of convergence.
[16]Much less for the *Stateless* algorithms, but in the order of millions of iterations for the *Stateful*.

revenues, individual bids, and rewards (vertical axis) by the percent of 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 *high-value* player due to the increase in the bid of the emphmiddle-value player.

Drivers of the Baseline Findings

Most AIAs lack the transparency and interpretability that would be required to explain the mechanism behind their decisions. The Q-learning algorithm, however, allows us to explore these mechanisms since the learning occurs through a well-defined "temporal difference" updating rule. Still, because of multi-agent simultaneous learning, the process is nonstationary and there is no available theoretical framework to formalize it. Therefore, the analysis of the drivers is a challenging task and it is in itself a relevant contribution from our study.[17]

The behavior of the *middle-value* player turns out to be fundamental to pin down the differences in revenues between *Stateful Synchronous* and *Stateless Asynchronous* algorithms. In *Stateful Synchronous* experiments, the *middle-value* player lowers her bid to 1.2 in order to minimize her price per click which is defined by the bid of the lowest player.[18] As soon as we force the *middle-value* player to

[17]In an ongoing follow-up study, (Decarolis et al. 2023), we propose methods to evaluate the learning behavior of Q-learning algos in a more general setting and explore implications for competition policy.

[18]We have confirmed the robustness of these findings by considering the finer grid, in which case we found that the *middle-value* player converges to the bid that is one step higher than 1 which is the valuation of the lowest bidder.

consistently bid anything higher than 1.2, *low-value* player expands the range of bids played which goes from the lowest possible bid and to the bid below the one of *middle-value* player. That necessarily leads to a higher price paid by the *middle-value* player. The main driving force of this result is the possibility of calculating the counterfactual rewards, in particular the *Synchronous* update. While learning, the lowest player at each iteration updates the Q-matrix for the current state for each of the actions. As a result, all bids below the bid of the *middle-value* player lead to the same current reward. And that forces the *middle-value* player to have to keep bidding low in order to minimize the price paid. Moreover, the dynamic strategies are also crucial for learning, since they allow the *middle-value* player to learn by associating specific actions to the higher future rewards. The logic can also be seen clearly in panel (c) of Figure 2. Here, we have focused on the *Stateful Synchronous* algorithm's bids at convergence. We start with the final Q-matrix, but then introduce an exogenous shock to the bid of *middle-value* player. Instead of bidding 1.2, she deviates to bidding 1.6 in the period that we call 0. What follows is that when *middle-value* player deviates, *low-value* player increases her bid (responds with a higher price).[19] Importantly, just after a few periods, all players return to equilibrium bidding.[20]

Instead, and surprisingly, when the *Stateless Asynchronous* algorithms are at play, there is no need for the *middle-value* player to lower her bid in order to force the *low-value* player not to play 1.2 or 1.4 when *middle-value* player converges to 1.6. The main idea here is that there is no *Synchronous* update in this case, as well as there is no memory of the past. Q-values only depend on the action since there is only one state of the world, and Q-values of the *low-value* player for all her actions at convergence are equal to zero. Instead, in a context where exploration is in place for several periods, the Q-values of actions above 1 converge to negative values in the long run. Indeed, those actions expose the *low-value* player to the risk

[19] This comes from a long history of learning, not a random action chosen by *low-value* player on a wider interval of actions as the Q-value of the bid 1.2 could have been updated just by an insignificant amount.

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

of being unprofitably undercut by random deviations of the opponents. Because her Q-values are generally small in absolute terms, negative rewards are heavily weighted in the updating process and she quickly learns that she should refrain from playing above her valuation. Since there is no *Synchronous* update in that case there is no possibility for *middle-value* player to learn that her other possible actions above the value of the *low-value* player lead to the same immediate rewards after *middle-value* player stabilizes at 1.6 and *low-value* player iterates between actions from 0.2 to 1.

To illustrate the importance of dynamic strategies in the *Stateful Synchronous* case, and building upon our earlier discussion of the experiments in Table 1, we compare the bid evolution when the synchronous algorithm learns to bid with different discounting of the future payoffs. In particular, panels (a) and (b) of Figure 2 show the evolution of the bids by the percent of 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.

VI. Generalizations

In this section, we extend our results along three dimensions. First, we consider alternative settings of Q-learning AIAs; 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. These results are presented in Appendix C. In particular, we first consider two alternative bid selection methods for the cases in which different actions are associated with identical Q-values. We show that the main results also hold both when the algorithms are *conservative* and choose the lowest bid, as well as when the algorithms are *greedy* and choose the highest bid. See Table A1.[21] In another extension, we consider asymmetric grids, spanning the same bid space, but featuring a different number of actions for each player. In particular, in Table A2 we show that results are almost identical for a variation of the baseline setting such that *high-value* player and *low-value* player have $k = 15$ possible bids, whereas *middle-value* player instead has a finer grid with $k = 20$ bids.[22]

The second generalization 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 82%.

The third and last variation involves alternative auction designs. In Appendix E, we consider another mechanism, namely the Vickrey-Clarke-Groves (VCG) which is also used in online advertising.[23] 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.[24]

[21]As a reminder, in the baseline model, we employ a *random* method - i.e., the AIA chooses the action randomly among all that lead to the highest possible Q-value in a given state.

[22]In that setting, the ties with *middle-value* player are no longer possible for all the bids excluding the $mix = 0.2$ and $max = 3$.

[23]VCG is allegedly used by Facebook.

[24]For instance, a thorough analysis of first price vs second price auctions is conducted in (Banchio and Skrzypacz 2022). This study considers exclusively stateless algorithms and focuses on a setting with two

VII. Conclusions

Our results highlight several important features regarding how AIAs shape the functioning of ad auctions, as the AIAs produce different outcomes relative to what might be expected under an equilibrium analysis. First, bids converge to values that are lower than the EOS equilibrium generally adopted in the literature. Second, the auctioneer prefers to provide bidders with less information. Third, the auctioneer revenues differ between GSP and VCG auctions, with the latter being lower.

The results of this study open the door to further extensions. The first avenue of research regards exploring the potential risk of abuses by platforms in terms of data usage. For instance, the deployment of proprietary AI bidding tools by the platform (like Google's Performance Max) poses a risk that the training occurs on data that is superior to that available to advertisers or their intermediaries, thus leading such platform-sponsored tools to outperform competition. A second direction involves extending the analysis of AIAs to a broader class of auction mechanisms. Especially in environments with asymmetric bidders and multiple, heterogenous items for sale little is known about the ranking of different auction formats, but the methodological approach proposed in this study offers a way to achieve that. Third, the understanding of the mechanisms through which AIAs bid could be expanded to guide data policy interventions and statistical testing into what might constitute collusive bidding.

symmetric players, asking whether bids converge to the Nash equilibrium. It finds that this is the case for the second-price auction, but bids are below the competitive level for the first-price auction. However, even in the first price auction bids converge to Nash equilibrium if information about the lowest bid to win is available. This indicates that stateless synchronous algorithms converge to Nash equilibria. A result which is also coherent with (Asker, Fershtman and Pakes 2022) who find the same in the context of a Bertrand oligopoly pricing game.

REFERENCES

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature*, 54(2): 442–92.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2022. *Prediction Machines, Updated and Expanded: The Simple Economics of Artificial Intelligence.* Harvard Business Review Press, Cambridge.

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

Asker, John, Chaim Fershtman, and Ariel Pakes. 2022. "Artificial intelligence, algorithm design, and pricing." *AEA Papers and Proceedings*, 112: 452–56.

Asker, John, Chaim Fershtman, and Ariel Pakes. 2023. "The Impact of AI Design on Pricing." *Journal of Economics and Management Strategies*, forthcoming.

Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu. 2023. "Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market." *Journal of Political Economy*, forthcoming.

Athey, Susan, and Denis Nekipelov. 2014. "A Structural Model of Sponsored Search Advertising Auctions." *Working paper.*

Banchio, M., and A. Skrzypacz. 2022. "Artificial Intelligence and Auction Design." *Working Paper.*

Banchio, M., and G. Mantegazza. 2022. "Adaptive Algorithms and Collusion via Coupling." *Working Paper.*

Bergemann, Dirk, and Alessandro Bonatti. 2019. "Markets for Information: An Introduction." *Annual Review of Economics*, 11(1): 85–107.

Bergemann, Dirk, and Stephen Morris. 2019. "Information Design: A Unified Perspective." *Journal of Economic Literature*, 57(1): 44–95.

Bergemann, Dirk, Paul Duetting, Renato Paes Leme, and Song Zuo. 2022*a*. "Calibrated Click-Through Auctions." *ACM WWW Conference Proceed-*

ings.

Bergemann, Dirk, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter. 2022*b*. "Optimal Information Disclosure in Classic Auctions." *American Economic Review: Insights*, 4(3): 371–88.

Bichler, Martin, M. Fichtl, and Stefan Heidekrüger. 2021. "Learning equilibria in symmetric auction games using artificial neural networks." *Nature Machine Intelligence*, 3: 687–695.

Blake, Thomas, Chris Nosko, and Steven Tadelis. 2015. "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment." *Econometrica*, 83(1): 155–174.

Borgers, Tilman, Ingemar Cox, Martin Pesendorfer, and Vaclav Petricek. 2013. "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence." *American Economic Journal: Microeconomics*, 5(4): 163–187.

Brown, Zach Y, and Alexander MacKay. 2023. "Competition in pricing algorithms." *American Economic Journal: Microeconomics*, 15(2): 109–156.

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

Chen, Le, Alan Mislove, and Christo Wilson. 2016. "An empirical analysis of algorithmic pricing on amazon marketplace." *Proceedings of the 25th international conference on World Wide Web*, 1339–1349.

Che, Yeon-Koo, Syngjoo Choi, and Jinwoo Kim. 2017. "An experimental study of sponsored-search auctions." *Games and Economic Behavior*, 102: 20 – 43.

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, 111(10): 3299–3327.

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, Tommaso Pellegrinetti, Gabriele Rovigatti, Michele Rovigatti, and Ksenia Shakhgildyan. 2023. "Multi-agent learning in auctions." *Working Paper.*

Deng, Yuan, Negin Golrezaei, Patrick Jaillet, Jason Cheuk Nam Liang, and Vahab Mirrokni. 2023. "Multi-channel Autobidding with Budget and ROI Constraints." *arXiv preprint arXiv:2302.01523.*

Dütting, Paul, Zhe Feng, Harikrishna Narasimhan, and David C. Parkes. 2017. "Optimal Auctions through Deep Learning." *CoRR*, abs/1706.03459.

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.

Fainmesser, Itay P, Andrea Galeotti, and Ruslan Momot. 2023. "Digital privacy." *Management Science*, 69(6): 3157–3173.

Goldfarb, Avi, and Verina F Que. 2023. "The Economics of Digital Privacy." *Annual Review of Economics*, 15.

Gomes, Renato, and Kane Sweeney. 2014. "Bayes-Nash Equilibria of the Generalized Second Price Auction." *Games and Economic Behavior*, 86: 421–437.

Heidekrüger, Stefan, Paul Sutterer, Nils Kohring, Maximilian Fichtl, and Martin Bichler. 2021. "Equilibrium Learning in Combinatorial Auctions: Computing Approximate Bayesian Nash Equilibria via Pseudogradient Dynamics." *CoRR*, abs/2101.11946.

IAB. 2018. "Artificial Intelligence: Myth versus reality in the digital advertising world." *Xaxis and Internet Advertising Board Report.*

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

Lahaie, Sébastien. 2006. "An analysis of alternative slot auction designs for sponsored search." *Proceedings of the 7th ACM Conference on Electronic Commerce*, 218–227.

Lambin, Xavier. 2023. "Less than meets the eye: simultaneous experiments as a source of algorithmic seeming collusion." *Available at SSRN 4498926*.

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 multi-player games with auction applications." *Working paper*.

Morton, Fiona M Scott, and David C Dinielli. 2020. "Roadmap for a digital advertising monopolization case against google." *Omidyar Network*, 1–38.

Simonov, Andrey, and Shawndra Hill. 2021. "Competitive advertising on brand search: Traffic stealing and click quality." *Marketing Science*, 40(5): 923–945.

Simonov, Andrey, Chris Nosko, and Justin M Rao. 2018. "Competition and crowd-out for brand keywords in sponsored search." *Marketing Science*, 37(2): 200–215.

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—: Simulation Results of Baseline Experimental Designs

	Bids	Individual Rewards	Revenue
Stateful Synchronous (Full Info)	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2
			[7.03, 7.37]
Stateless Asynchronous (Limited Info)	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76
			[8.39, 9.13]
Partial Asynchronous	$(2.2, 1.36, 0.59)$	$(7.74, 2.62, 0.13)$	7.87
			[7.35, 8.39]
Stateful Synchronous ($\delta = 0$)	$(2.46, 1.47, 0.61)$	$(7.64, 2.79, 0.0)$	8.57
			[8.24, 8.9]
Stateless Synchronous	$(2.49, 1.49, 0.6)$	$(7.55, 2.8, 0.0)$	8.65
			[8.31, 8.99]

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 experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless experiments were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

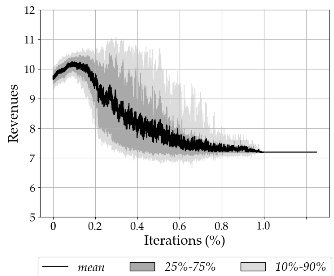

(a) Stateful Sync Revenues

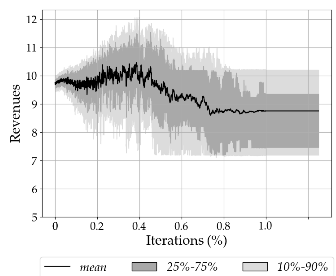

(b) Stateless Async Revenues

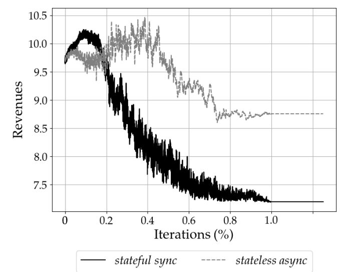

(c) Revenues Evolutions

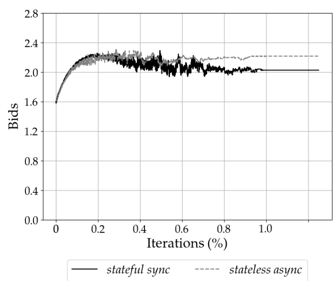

(d) Bids of Player $v_1 = 3$

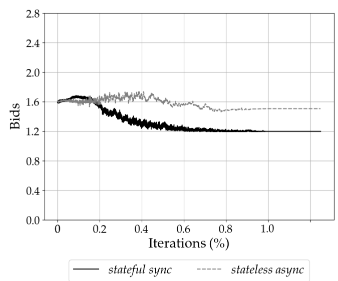

(e) Bids of Player $v_2 = 2$

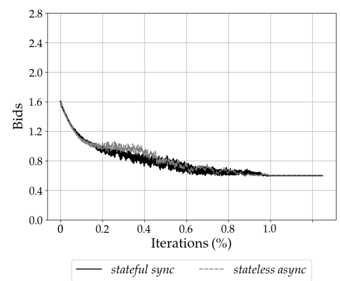

(f) Bids of Player $v_3 = 1$

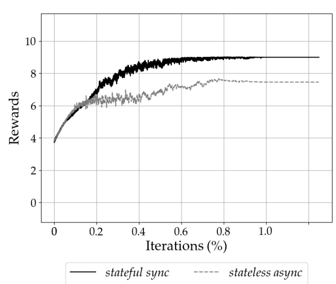

(g) Rewards of Player $v_1 = 3$

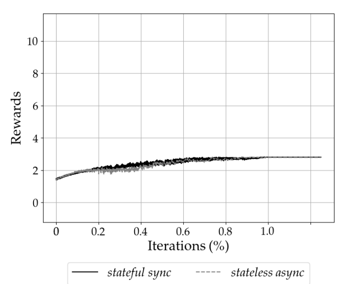

(h) Rewards of Player $v_2 = 2$

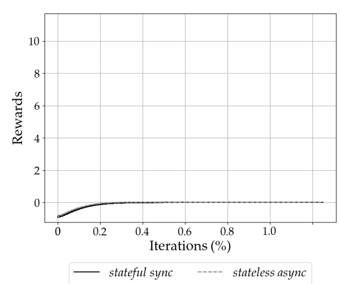

(i) Rewards of Player $v_3 = 1$

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 of 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 of the total number of iterations (horizontal axis). The Stateful Synchronous experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e-07$, $\delta = 0.95$. The Stateless Asynchronous experiments, were run with 1,000,000 iterations and $\beta = 1.22e-05$.

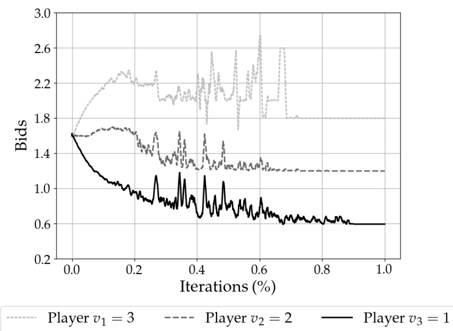

(a) Bids when $\delta = 0.95$

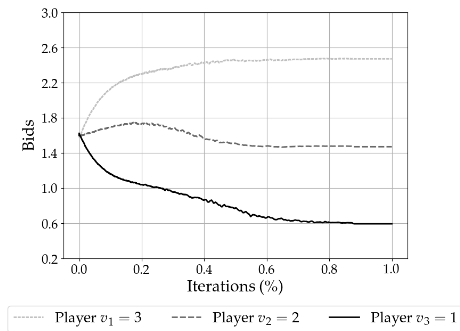

(b) Bids when $\delta = 0$

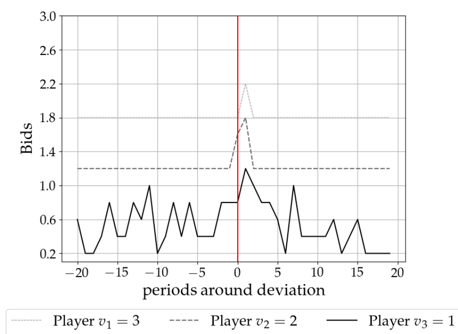

(c) Player $v_2 = 2$ Deviates

Figure 2. : Evolution of Bids in a Single Run of Stateful Synchronous Experiment

Note: Panel (a) shows the evolution of the smoothed bids (moving average was applied) by the percent of 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 of the total number of iterations in just one run of Stateful Synchronous experiment with $\delta = 0$. Panel (c) shows the evolution of the bids by the iteration from the moment of the forced deviation of *middle-value* player to raise her bid to 1.6 instead of the bid 1.2 at convergence of the Stateful Synchronous algorithm.Stateful Synchronous experiment was run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e-07$, and both $\delta = 0.95$, and $\delta = 0$.

Appendix

A. Data Policies

In this Appendix we discuss a few examples of data policies and their evolution. We focus on the case of Google. (Varian 2007) and (Edelman, Ostrovsky and Schwarz 2007) pioneered the equilibrium analysis of the search auctions and argued 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 estimate 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 forward to today, the situation is radically changed. 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. In particular, 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 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 to advertisers.[1] 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 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).[2] Several other examples exist and all describe the same pattern toward data obfuscation.[3]

Therefore, 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 consider (*Limited Information*) is likely a reasonable approximation of how this market currently works. Indeed, as much as it appears ludicrous that an advertiser does not even observe the price she 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 is assigned 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.[4]

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

[1] 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.

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

[3] 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.

[4] Attribution is a fundamental element in digital advertising, as 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 cryptic description: "data-driven attribution uses the Shapley value solution concept from cooperative game theory". See details in Figure A2 in the Online Appendix.

to data policies adopted by the large digital platforms.[5]

B. Convergence

Since the rewards, as well as the update rule in multi-agent reinforcement learning generally depend on the actions of other players, each player's optimization problem is inherently non-stationary. For example, the same chosen action can lead to different rewards based on the bids of other players. When the bidder's competitors change actions over time, because of experimentation and learning, the optimization problem of each bidder becomes nonstationary. Such nonstationarity is at the root of the lack of general convergence results for Q-learning in games. Moreover, even if the algorithms converge, the basic question is whether the bids converge toward the Nash Equilibrium (NE) predictions. Importantly, 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 and $\beta = 1.22e - 05$, were enough for the convergence to be reached with enough exploration. That is because the Q-matrix in this case takes the form of a Q-vector since the state is a singleton. The convergence check covered the last 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 5e-06, so that convergence can be reached:

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

$exp(-15,000,000 * 8.137e-07) = 5e-06$. 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 experiments with 3 players (each cell is visited more than 100 times). As can be seen in Table A3, each cell of the Q-matrix is visited at least 117 times. In some simulations, the algorithms experience cyclical behavior and do not bid a constant amount.

C. Alternative Q-learning AIAs Settings

In this appendix, we explore alternative settings of Q-learning AIAs. We first consider two alternative bid selection methods for the cases in which different actions are associated with identical Q-values. That is especially often the case under the *Sychronous Updating* since in the GSP auction several bids lead to the same reward given the bids of other players if the bidder's position stays the same. As a reminder, in the baseline model, we employ a *random* method - i.e., the AIA chooses the action randomly among all that lead to the highest possible Q-value in a given state. We show that the main results also hold both when the algorithms are *conservative* and choose the lowest bid, as well as when the algorithms are *greedy* and choose the highest bid. In Table A1, we present results similar to the ones in Table 1 under different bid selection methods.

In the case of conservative algorithms, 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. In the case of greedy algorithms, the decision not to reveal competitor bids increases the platform's average revenues by 25% from 8 to 10. That is driven by the reduction in the reward of the highest-value player from 9 to 7 due to the increase in the bid of the second-highest-value player, from 1.2 to 1.6. In sum, the only difference with the baseline setting is the bidding behavior of *low-value* player. If in the baseline setting, we observe the lowest bidder to submit at convergence different bids between 0.2 and 1, in the case of *conservative* setting, she converges to bidding 0.2, whereas in case of *greedy* algorithms, to 1.

In another extension, we consider asymmetric grids, spanning the same bid space, but featuring a different number of actions for each player. In particular, in Table A2 we show that results are almost identical for a variation of the baseline setting such that *high-value* player and *low-value* player have $k = 15$ possible bids, whereas *middle-value* player instead has a finer grid with $k = 20$ bids. In that setting, the ties with Player v_2 are no longer possible for all the bids excluding the $mix = 0.2$ and $max = 3$. Thus, we have shown that for our results it is not crucial to have the same grids of actions for each of the players.

D. Milgrom and Mollner (2018) 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 $[B_{\min,} B_{\max}] = [1, 15]$, 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 82% increase. Thus, the increase in auctioneer revenues due to the information restriction can be much higher than in our baseline experiment.

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. Compared to the 22%

increase for GSP, 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 setting.

Table A1—: Simulation Results in Case of Conservative and Greedy AIAs

	Bids	Individual Rewards	Revenue
Stateful Synchronous (Conservative)	$(2.05, 1.2, 0.2)$	$(9.0, 3.6, 0.0)$	6.4 [6.4, 6.4]
Stateless Asynchronous (Conservative)	$(2.14, 1.49, 0.2)$	$(7.49, 3.59, 0.0)$	7.86 [7.56, 8.16]
Stateful Synchronous (Greedy)	$(2.0, 1.2, 1)$	$(9.0, 2.0, 0.0)$	8.0 [8.0, 8.0]
Stateless Asynchronous (Greedy)	$(2.21, 1.6, 1.0)$	$(7.0, 2.0, 0.0)$	10.0 [9.67, 10.33]

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, or greedy and choose the biggest 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 Synchronous experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Asynchronous experiments, were run with 1,000,000 iterations and $\beta = 1.22e-05$.

Table A2—: Simulation Results in Case when *middle-value* player has a Grid of 20 Bids

	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(2.03, 1.23, 0.66)$	$(8.84, 2.67, 0.0)$	7.49 [7.2, 7.78]
Stateless Asynchronous	$(2.17, 1.45, 0.61)$	$(7.73, 2.78, 0.0)$	8.49 [8.06, 8.91]

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 Synchronous experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Asynchronous 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 Updates	3375.0	64369.79	515598.71	1769.1	2316.3	4715.7	11493.3	14320122.6
Cell Updates	50625.0	4291.32	34368.49	117.94	154.4	314.38	766.48	954674.84

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. 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.

Table A4—: Simulation Results in Alternative Example from Milgrom and Mollner (2018)

	Bids	Individual Rewards	Revenue
Stateful Synchronous	$(11.45, 4.22, 2.12)$	$(1078.0, 23.63, 5.0)$	428.36
			[416.68, 440.05]
Stateless Asynchronous	$(13.36, 7.72, 2.01)$	$(728.0, 23.96, 5.0)$	778.04
			[711.29, 844.8]

Note: In this example, $v_1 = 15$, $v_2 = 10$, $v_3 = 5$,while the click-through rates are $x^1 = 100$, $x^2 = 3$ and $x^3 = 1$. 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 Synchronous experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Asynchronous experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

Table A5—: Comparison of the GSP and VCG

	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	$(2.53, 1.21, 0.6)$	$(10.18, 2.8, 0.0)$	6.02
			[7.03, 7.37]			[5.66, 6.37]
Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76	$(2.79, 1.94, 0.62)$	$(7.94, 2.76, 0.0)$	8.3
			[8.39, 9.13]			[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 Synchronous experiments were run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, $\delta = 0.95$. The Stateless Asynchronous experiments, were run with 1,000,000 iterations and $\beta = 1.22e - 05$.

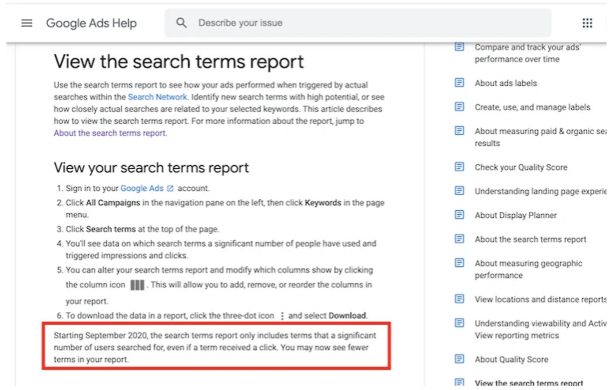

(a) Search Term Report

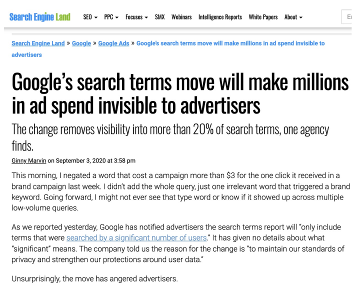

Screenshot taken from URL https://searchengineland.com/googles-search-terms-move-will-make-millions-in-ad-spend-invisible-to-advertisers-340182 on Feb. 20, 2023

(b) Impacts of the Search Term Report Change

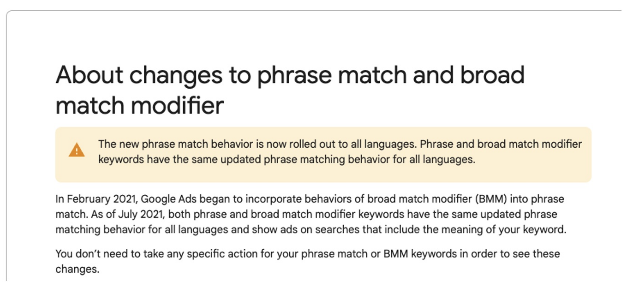

Screenshot taken from URL https://support.google.com/google-ads/answer/10286719?hl=en on Feb. 20, 2023

(c) Broad Match Modifiers

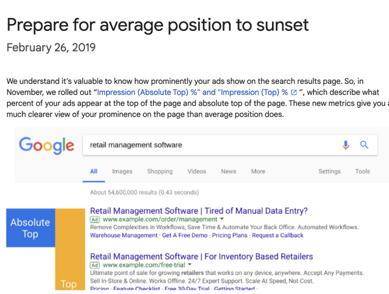

Screenshot taken from URL https://support.google.com/google-ads/answer/9263492 on Feb. 20, 2023

(d) Position Information

Figure A1. : Examples of Data Policy Changes

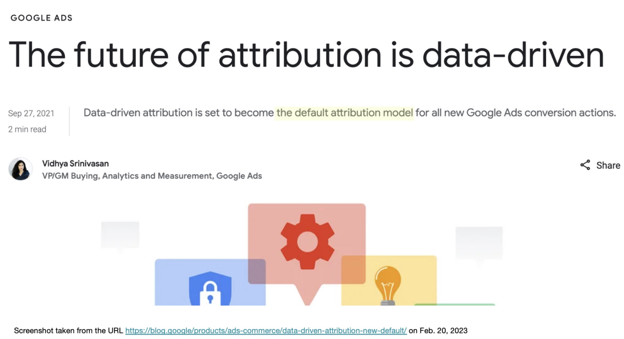
GOOGLE ADS

The future of attribution is data-driven

Sep 27, 2021
2 min read

Data-driven attribution is set to become the default attribution model for all new Google Ads conversion actions.

Vidhya Srinivasan
VP/GM Buying, Analytics and Measurement, Google Ads

Share

Screenshot taken from the URL https://blog.google/products/ads-commerce/data-driven-attribution-new-default/ on Feb. 20, 2023

(a) Attribution

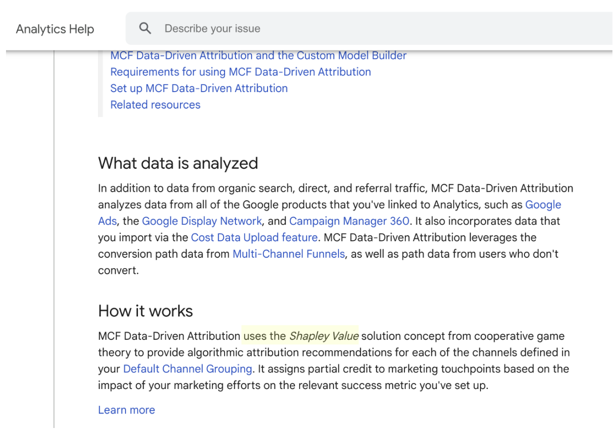
Analytics Help

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