

DISCUSSION PAPER SERIES

DP18009

**ARTIFICIAL INTELLIGENCE & DATA
OBFUSCATION: ALGORITHMIC
COMPETITION IN DIGITAL AD
AUCTIONS**

Francesco Decarolis, Gabriele Rovigatti, Michele
Rovigatti and Ksenia Shakhgildyan

INDUSTRIAL ORGANIZATION

CEPR

ARTIFICIAL INTELLIGENCE & DATA OBFUSCATION: ALGORITHMIC COMPETITION IN DIGITAL AD AUCTIONS

Francesco Decarolis, Gabriele Rovigatti, Michele Rovigatti and Ksenia Shakhgildyan

Discussion Paper DP18009

Published 20 March 2023

Submitted 20 March 2023

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Industrial Organization

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Francesco Decarolis, Gabriele Rovigatti, Michele Rovigatti and Ksenia Shakhgildyan

ARTIFICIAL INTELLIGENCE & DATA OBFUSCATION: ALGORITHMIC COMPETITION IN DIGITAL AD AUCTIONS

Abstract

Data are the key fuel of artificial intelligence and any change to the type and quality of available data has an impact on the type and performance of the feasible algorithms. We analyze the incentives that large digital platforms have to alter data flows to their advantage by strategically obfuscating data. We quantify this phenomenon in the context of digital advertising auctions through a series of simulated experiments where asymmetric bidders employ artificial intelligence algorithms to compete in Generalized second-price auctions. We find that when less detailed information is available to train algorithms, auctioneer revenues are substantially and persistently higher.

JEL Classification: C73, D82, D83, D18

Keywords: Auctions, Procurement, Collusion, Data, Privacy

Francesco Decarolis - francesco.decarolis@unibocconi.it
Bocconi University and CEPR

Gabriele Rovigatti - gabriele.rovigatti@gmail.com
Bank of Italy

Michele Rovigatti - mic.rovigatti@gmail.com
Bocconi University

Ksenia Shakhgildyan - ksenia.shakhgildyan@unibocconi.it
Bocconi University

Acknowledgements

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

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

By FRANCESCO DECAROLIS, GABRIELE ROVIGATTI, MICHELE ROVIGATTI,
KSENIA SHAKHGILDYAN*

Data are the key fuel of artificial intelligence and any change to the type and quality of available data has an impact on the type and performance of the feasible algorithms. We analyze the incentives that large digital platforms have to alter data flows to their advantage by strategically obfuscating data. We quantify this phenomenon in the context of digital advertising auctions through a series of simulated experiments where asymmetric bidders employ artificial intelligence algorithms to compete in Generalized second-price auctions. We find that when less detailed information is available to train algorithms, auctioneer revenues are substantially and persistently higher.

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 in data flows motivated by various initiatives for privacy protection enacted by both governments and large digital platforms. For instance, through regulations like the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the US, governments are intervening in digital markets to establish and protect a right to privacy in the digital realm. The large digital platforms have similarly embraced privacy-motivated initiatives, like the tracking restrictions that Apple introduced with the launch of IOS14 in April 2021 or the Google announcement in March 2021 that it is going to eliminate third-party cookies. This type of actions, while motivated by consumers’ privacy concerns, crucially impact platform-to-business interactions. Although the ongoing debate about regulating digital platforms is looking at both these privacy initiatives and the growing use of AIAs, it is doing so as if they were separate phenomena. 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.

A growing literature in economics focuses on data, data markets, and data governance (see the reviews in (Bergemann and Morris 2019, Bergemann and Bonatti 2019)). However, the link between platform data strategies and the AIAs used by the business 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 extensively 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 artificial intelligence.¹

Specifically, we focus on sponsored search which is the most lucrative portion of digital advertising (worth more than 40 percent of all digital ad revenues).² 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 benefit from obfuscating the data available to the AIAs bidding in its auctions. 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 Digital Markets Act (DMA) in Europe mandate that ad platforms disclose data to advertisers, 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 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.⁴ In our simulated experiments, asymmetric bidders employ AIAs to compete in repeated 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, effectively, data policies determine to a significant extent the types of AIAs can be used.

We focus on the case of Q-learning algorithms to take advantage of the trans-

¹For an early industry survey, see (IAB 2018).

²See 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). Importantly, the Google's top revenue source is search ads: according to Statista, of the \$279.81bn in revenue in 2022, \$162.45bn (58%) came from search ads.

³The DMA contains provisions that on the one hand enhance transparency toward advertisers (Art. 6g), while on the other hand place restrictions on targeted and micro-targeted ads (Art. 6aa).

⁴In (Calvano et al. 2020) AIAs set prices in a Bertrand oligopoly, while we focus on auctions.

parency of the learning process entailed by these algorithms.⁵ We illustrate 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. We show that when we move between a scenario of full 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%. We also show that an intermediate case where some, partial information about rival bids is disclosed, leads to a more moderate decline in advertiser revenues.⁶

We then explore the mechanisms behind our 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 AIAs can be traced back to dynamic strategies.⁷ Through a series of controlled alterations of the baseline algorithms, we detect that it is the feasibility of dynamic strategies entailing the possibility of punishment-reward schemes that drive the tendency of the most data-intensive algorithms to coordinate on lower bids. Indeed, our second main finding is that algorithmic bidding has a tendency to sustain low bids relative to the competitive benchmark case and that this tendency is stronger the more data are available.

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

⁵Q-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.

⁶Note that we focus on the one keyword ad auction to show that already in this setting the effect of information obfuscation on auctioneer revenues is strong. In this setting, the crucial information regards the competitors' bids. In reality, since advertisers bid on many keywords, and there could exist additional uncertainty regarding the profitability of bidding on a given keyword, the overall effect would be stronger.

⁷For the practical relevance and fundamental role of dynamic strategies in repeated auctions see (Chassang and Ortner 2018).

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. We briefly list these literatures here, while leaving to later sections the discussion of how our results more directly relate to other studies. The first branch regards data markets and the design of information. While the relevance of information is an old idea, its modern incarnation in the form of studies on privacy, information design, and information markets reflects the key role that data has come to play in today’s digital economies. Three comprehensive overviews of these areas are (Acquisti, Taylor and Wagman 2016) about privacy, (Bergemann and Morris 2019) about information design and (Bergemann and Bonatti 2019) about information markets. 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. 2022b) and (Bergemann et al. 2022a). Our study contributes to this literature by proposing a new viewpoint to interpret platform data strategies: that of influencing the outcomes of the AIAs operating through platform-provided data. Although not focused on AI, the closest paper to ours is (Alcobendas, Kobayashi and Shum 2021) on how the elimination of third-party cookies (as announced by Google and already implemented by other browsers), by depriving advertisers of this information, affects their bidding and revenues in the display ad auctions. Similarly to us, they find marked losses for advertisers and benefits for the platform.

The second branch is that on the effects of artificial intelligence and, more specifically, its impacts on pricing. (Agrawal, Gans and Goldfarb 2022) offers a general overview of the economic impacts of AI, while (Calvano et al. 2020) is the first study to explore how and why prices set by AIAs might differ from static equilibrium predictions. How AIAs design features drive pricing has been stressed by (Asker, Fershtman and Pakes 2022). 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. 2020) who use data on gas pricing in Germany to show how the deployment of AIAs in this market lead to substantial price increases. This literature is rapidly growing in many directions,⁸ 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). Compared to these studies, our contribution differs both because we analyze AIAs bidding in GSP and VCG auctions (i.e., solving the problem of simultaneously allocating multiple, heterogeneous objects) and because we emphasize how the designs of the AIAs is reflective of the type of data available for their training.

The third branch is that on online advertising and, more specifically, on sponsored search. This key area of digital advertising has been the focus of intensive research and we are particularly close to those studies centered on the functioning of search auctions. 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 and the forthcoming regulation of this sector, (Morton and Dinielli 2020). In this respect, our results drive attention to a series of changes in the data flows characterizing this sector and offer a rationale for why they might concern regulators.

The paper proceeds as follows: Section 2 presents data policies and obfuscation strategies by the platform; Section 3 presents the GSP auction model; Section 4 describes the design features of the AIAs experiments; Section 5 describes the simulation results; Section 6 presents the generalizations; Section 7 concludes.

⁸See, among others, (Chen, Mislove and Wilson 2016), (Competition and Authority 2018), (Klein 2021), (Brown and MacKay 2021), (Banchio and Mantegazza 2022) and (Mehta and Perlroth 2023).

II. Data Policy: Obfuscation Strategies by the Platform

We start by describing the data policies considered in this paper and by providing some evidence on the obfuscation strategies of the platforms running online ad auctions. By data obfuscation, we refer to the tendency of search engines to disclose increasingly coarser data to advertisers.

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, in this paper, 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 of *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.

We now turn to 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.⁹ 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).¹⁰ Several other examples exist and all describe the same pattern toward data obfuscation.¹¹

⁹See 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.

¹⁰See panel (A1c) in Figure A1 in the Online Appendix reporting a screenshot of the Google announcement.

¹¹For 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

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 (*No 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.¹²

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.¹³

Without questioning the underlying motivations, it is important to assess what the consequences, intended or not, are. This is what we do in the next sections by evaluating which AIAs are feasible and how they perform under different information regimes. But we first introduce our baseline setting of the GSP auction.

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

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.

¹²Attribution 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.

¹³Apple, 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>.

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

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.¹⁴ 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.¹⁵ In particular, the AIAs that we consider are Q-learning algorithms bidding against each other and learning simultaneously.¹⁶ Their training entails striking a bal-

¹⁴A 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=tW3BRM1d1c8>. 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 $i = 2, \dots, S$, $b_i = v_i - \frac{x^i}{x^i - 1}(v_i - b_{i+1})$.

¹⁵Which equals for player i , $E[\sum_{t=0}^{\infty} \delta^t r_t^i]$, where δ is the discount factor.

¹⁶For an overview of this type of AIAs see (Sutton and Andrew 2018).

ance 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 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.¹⁷ We use a declining with time exploration rate $\epsilon_t = e^{-\beta * t}$, where $\beta > 0$ is the annihilation coefficient.¹⁸

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

IV.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

¹⁷In 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 B.

¹⁸We 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 5e-06.

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:

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

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

Here α is the learning rate.¹⁹ 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 bid b , given the bids of other players b^{-i} :

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

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.²⁰ Hence, if bidder i observes the

¹⁹In 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.

²⁰A synchronous algorithm thus operates in a way similar to how economists approach equilibrium analysis. See (Asker, Fershtman and Pakes 2022) for a discussion.

bids of competitors b^{-i} , the calculation of his reward under any possible b^i holding fixed b^{-i} is trivial and, therefore, 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 .

IV.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. In particular, in case of *Full Information*, we define for each of the players $s_t = (b_{t-1}^i, b_{t-1}^{-i})$, whereas for the case of *Partial Information*, $s_t = (b_{t-1}^i, b_{t-1}^{\rho(i)+1})$.
- 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. But 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 *No Information*, the *Stateless Asynchronous* (Stateless algorithms with Asyn-

chronous updating rule) is the only possible form of AIA. *Partial Information* allows asynchronous AIAs to retain some, limited form of memory containing one’s own bids and prices paid (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.

V. 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 *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.²¹ 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

²¹The 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.

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 the case of *No 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). 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 prices paid.

Two other cases are considered and reported in the last two rows of Table 2. 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 $\delta=0$ 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 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.

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.²² But convergence requires a large number of periods.²³ Figure 1 shows the evolution of the auctioneer’s 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 highest-value player due to the increase in the bid of the second-highest-value player.

V.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 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

²²See Appendix A for the definition and discussion of convergence.

²³Much less for the *Stateless* algorithms, but in the order of millions of iterations for the *Stateful*.

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 with 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).²⁴ Importantly, just after a few periods, all players return to equilibrium bidding.²⁵

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 B. 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 A1.²⁶ 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

²⁴That 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.

²⁵The exact number of periods depends on a particular simulation run, and varies from 2 to 5.

²⁶As 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.

such that Player $v_1 = 3$ and $v_3 = 1$ have $k = 15$ possible bids, whereas Player $v_2 = 2$ instead has a finer grid with $k = 20$ bids.²⁷

The second generalization regards variations to the GSP stage game. In Appendix C, 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 D, we consider another mechanism, namely the Vickrey-Clarke-Groves (VCG) which is also used in online advertising.²⁸ 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.²⁹

VII. 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

²⁷In that setting, the ties with Player v_2 are no longer possible for all the bids excluding the $mix = 0.2$ and $max = 3$.

²⁸VCG is allegedly used by Facebook.

²⁹For 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 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 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 2021) who find the same in the context of a Bertrand oligopoly pricing game.

by AI across several stylized auction games characterized by different information levels. We find that when less detailed information is available to the algorithms, the advertisers' rewards tend to be lower, and conversely, the auctioneer revenues tend to increase. In particular, the decision not to reveal competitor bids increases the platform revenues by 22%. This increase is driven by the fact that with *No Information* (the bidder observes only her own reward) only *Stateless 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 might help to understand the implications of recent decisions by digital platforms. In the context of search advertising, advertisers on Google have experienced a reduction in the amount and 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. This result extends to the case of search advertising a series of other studies that have emphasized side effects of the public and private initiatives motivated by privacy protection concerns, see (Aridor, Che and Salz 2022) and (Lefrere et al. 2022).

Our result also highlights several important features regarding how AI shapes the functioning of ad auctions. In our setting, the AIAs produce three outcomes that differ relative to what might be expected under an equilibrium analysis. First, bids converge to values that are lower than those of the EOS equilibrium generally adopted in the literature. Second, the auctioneer revenues differ between GSP and VCG auctions, with the latter being lower. Third, contrary to the common wisdom in the auction literature (i.e., the linkage principle), for the auctioneer it is preferable to provide bidders with less information. The spirit of this third result is thus analogous to that of (Alcobendas, Kobayashi and Shum 2021) who

identify an incentive for a platform to obfuscate data in the form of elimination of third-party cookies in display ad auctions. In their case, the incentive for the platform derives from its maintained access to the cookie data that can be used to advantage selected bidders. In our case, what solves the puzzling contradiction with the linkage principle is that, when AIAs bid, their tendency toward lower bids is enhanced by access to more information, thus outweighing the benefits that an auctioneer would derive from competition among better informed bidders.

Overall, the results in this study open the door to numerous extensions. For instance, in a follow-up study, we also analyze the tension created by the deployment of AI bidding tools by the platform. These tools 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 for advertisers are significant.

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*.
- 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*.
- Athey, Susan, and Denis Nekipelov.** 2014. “A Structural Model of Sponsored Search Advertising Auctions.” *mimeo*.
- 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.** 2022a. “Calibrated Click-Through Auctions.” *ACM WWW Conference Proceedings*.
- Bergemann, Dirk, Tibor Heumann, Stephen Morris, Constantine Sorokin, and Eyal Winter.** 2022b. “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, 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.
- Chassang, Sylvain, and Juan Ortner.** 2018. “Collusion in Auctions with Constrained Bids: Theory and Evidence from Public Procurement.” *Journal of Political Economy*, Forthcoming.
- 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.

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.

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.
- 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 Perloth.** 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.” Citeseer.
- 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
<i>Stateful Synchronous</i>	(2.03, 1.2, 0.6)	(9.0, 2.8, 0.0)	7.2 [7.03, 7.37]
<i>Stateless Asynchronous</i>	(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 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 2—: Simulation Results of Alternative Experimental Designs

	Bids	Individual Rewards	Revenue
<i>Partial Asynchronous</i>	(2.2, 1.36, 0.59)	(7.74, 2.62, 0.13)	7.87 [7.35, 8.39]
<i>Stateful Synchronous</i> ($\delta = 0$)	(2.46, 1.47, 0.61)	(7.64, 2.79, 0.0)	8.57 [8.24, 8.9]
<i>Stateless Synchronous</i>	(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 auction 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 $\beta = 1.22e - 05$.

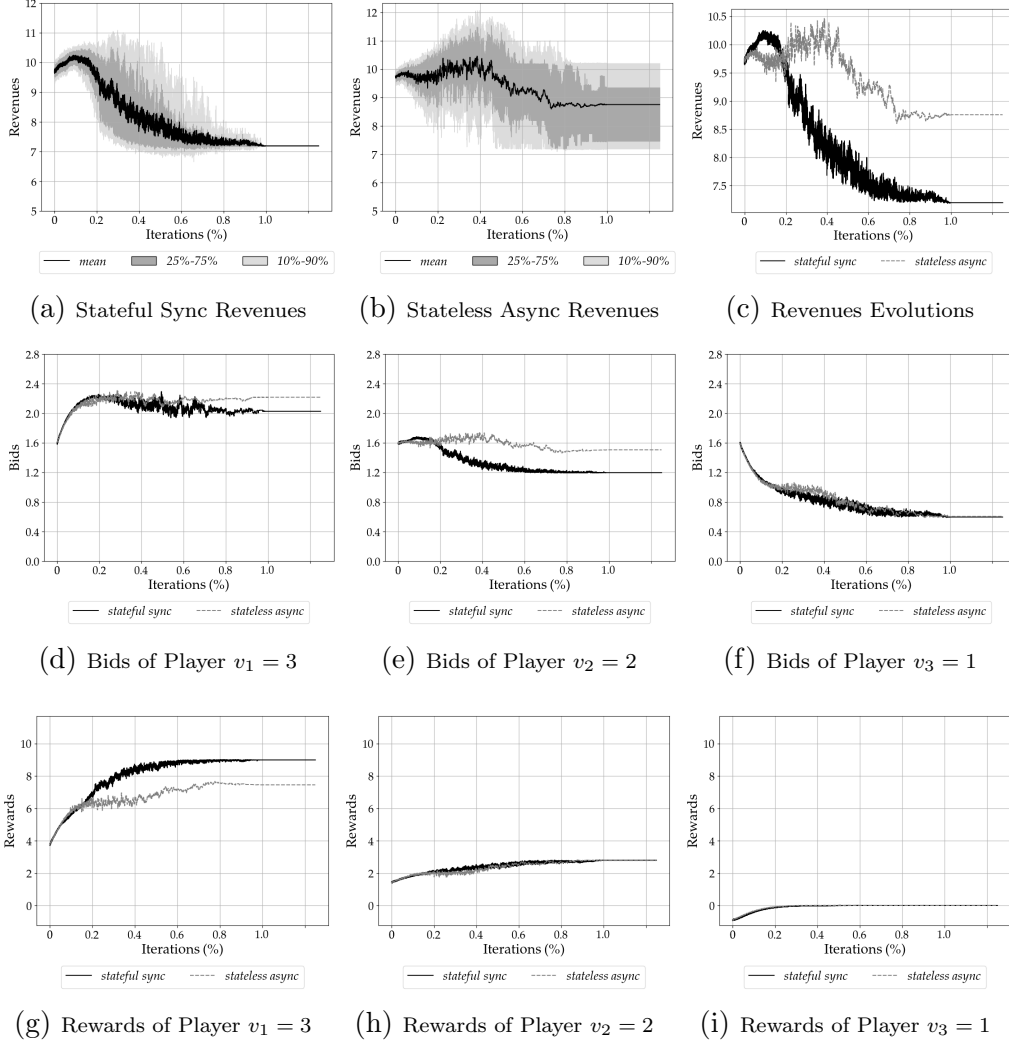


Figure 1. : Evolution of Auctioneer's Revenues, Bids, and Rewards

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 $\beta = 1.22e - 05$.

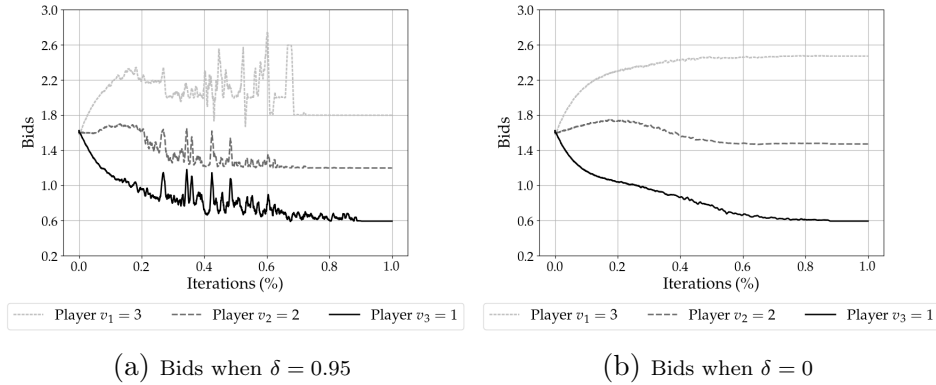


Figure 2. : Evolution of Bids in a Single Run of Stateful Synchronous Experiment

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 $\delta = 0$. Stateful Synchronous experiment was run with 15,000,000 iterations, $\alpha = 0.1$, $\beta = 8.137e - 07$, and both $\delta = 0.95$, and $\delta = 0$.

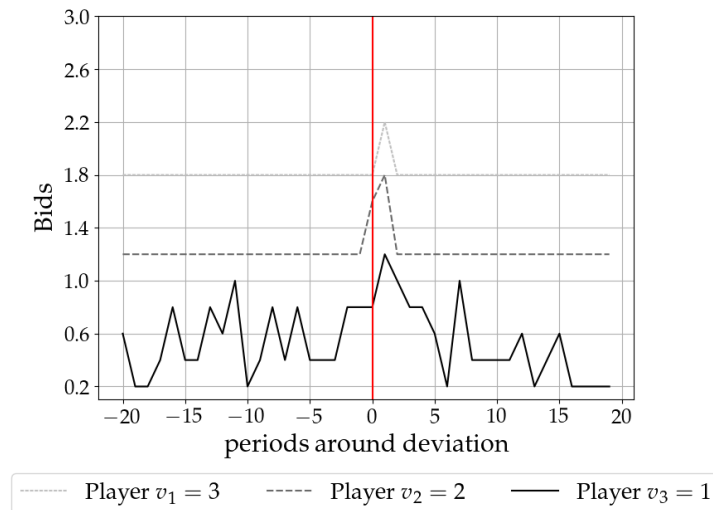


Figure 3. : Evolution of Bids in a Single Run of the Stateful Synchronous Algorithm when Player $v_2 = 2$ Deviates

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

A. 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 $b_{i,t}(s) = \operatorname{argmax}[Q_{i,t}(b, s)]$ 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: $\exp(-15,000,000 * 8.137e - 07) = 5e - 06$. If the rival is experimenting at even a 1% rate, the environment is still too non-stationary for the algorithm to con-

verge. 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.

B. 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 *Synchronous 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 Player $v_3 = 1$. 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 Player $v_1 = 3$ and $v_3 = 1$ have $k = 15$ possible bids, whereas Player $v_2 = 2$ 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.

C. 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 $[B_{\min}, B_{\max}] = [1, 15]$, 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.

D. 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
<i>Stateful Synchronous (Conservative)</i>	(2.05, 1.2, 0.2)	(9.0, 3.6, 0.0)	6.4 [6.4, 6.4]
<i>Stateless Asynchronous (Conservative)</i>	(2.14, 1.49, 0.2)	(7.49, 3.59, 0.0)	7.86 [7.56, 8.16]
<i>Stateful Synchronous (Greedy)</i>	(2.0, 1.2, 1)	(9.0, 2.0, 0.0)	8.0 [8.0, 8.0]
<i>Stateless Asynchronous (Greedy)</i>	(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 Player $v_2 = 2$ has a Grid of 20 Bids

	Bids	Individual Rewards	Revenue
<i>Stateful Synchronous</i>	(2.03, 1.23, 0.66)	(8.84, 2.67, 0.0)	7.49 [7.2, 7.78]
<i>Stateless Asynchronous</i>	(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 (2014)

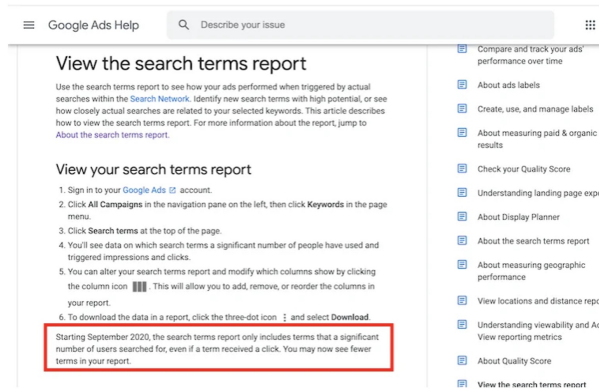
	Bids	Individual Rewards	Revenue
<i>Stateful Synchronous</i>	(11.45, 4.22, 2.12)	(1078.0, 23.63, 5.0)	428.36 [416.68, 440.05]
<i>Stateless Asynchronous</i>	(13.36, 7.72, 2.01)	(728.0, 23.96, 5.0)	778.04 [711.29, 844.8]

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 $\beta = 1.22e - 05$.

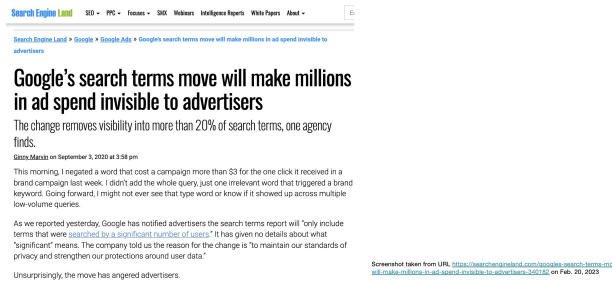
Table A5—: Comparison of the GSP and VCG

	GSP			VCG		
	Bids	Individual Rewards	Revenue	Bids	Individual Rewards	Revenue
<i>Stateful Synchronous</i>	(2.03, 1.2, 0.6)	(9.0, 2.8, 0.0)	7.2 [7.03, 7.37]	(2.53, 1.21, 0.6)	(10.18, 2.8, 0.0)	6.02 [5.66, 6.37]
<i>Stateless Asynchronous</i>	(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]

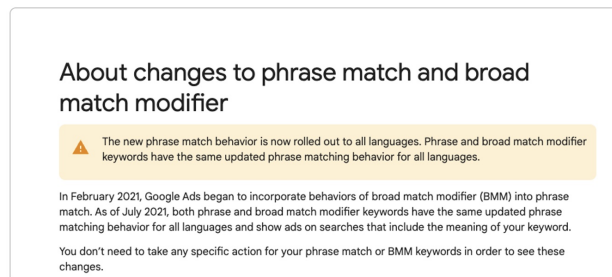
Note: Column 1 reports the average across runs limit bids for each of the players in the decreasing order of valuations; column 2 - the average across runs limit individual reward for each of the players in the decreasing order of valuations; while column 3 - the average auctioneer revenue across runs as well as the 95% confidence interval for the revenue in squared brackets. The Stateful 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



(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

Prepare for average position to sunset
February 26, 2019

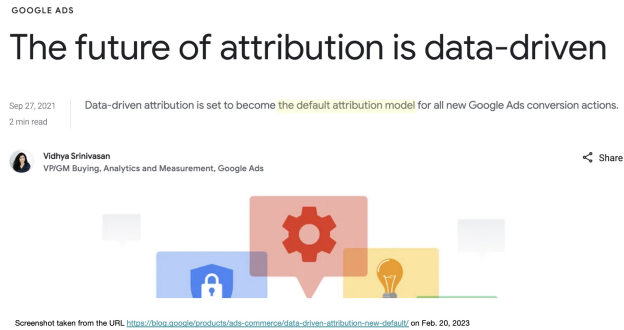
We understand it's valuable to know how prominently your ads show on the search results page. So, in November, we rolled out "Impression (Absolute Top) %" and "Impression (Top) % of 1", which describe what percent of your ads appear at the top of the page and absolute top of the page. These new metrics give you a much clearer view of your prominence on the page than average position does.



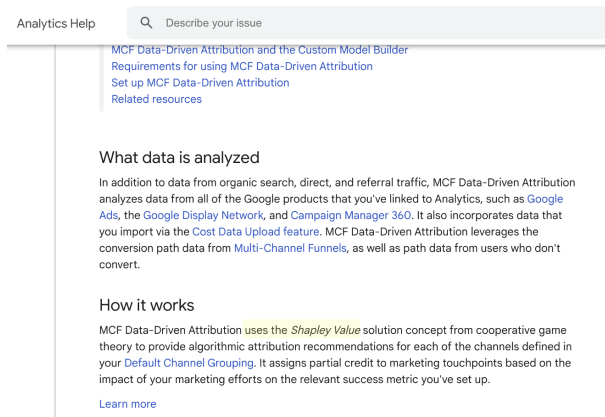
Screenshot taken from URL <https://support.google.com/google-ads/answer/9283492> on Feb. 20, 2023

(d) Position Information

Figure A1. : Examples of Data Policy Changes



(a) Attribution



(b) Details on the Attribution Model

Figure A2. : Changes in Default Attribution