AI & Data Obfuscation: Algorithmic Competition in Digital Ad Auctions

Francesco Decarolis (Bocconi University)
Gabriele Rovigatti (Bank of Italy)
Michele Rovigatti (Bocconi University)
Ksenia Shakhgildyan (Bocconi University)

EEA-ESEM 2023
Barcelona, August 30, 2023

Interaction between Algorithmic Pricing & Data Policies

Two interrelated trends characterize digital markets:

- AI algorithms (AIAs) are increasingly used for pricing and other economic decisions

Interaction between Algorithmic Pricing & Data Policies

Two interrelated trends characterize digital markets:

- AI algorithms (AIAs) are increasingly used for pricing and other economic decisions
- Public and private initiatives are changing data policies

Interaction between Algorithmic Pricing & Data Policies

Two interrelated trends characterize digital markets:

- AI algorithms (AIAs) are increasingly used for pricing and other economic decisions
- Public and private initiatives are changing data policies

Regulations in the EU and elsewhere are tackling both (especially the latter), but separately → potentially problematic since data is the AIAs' fuel!

Our research question: do digital platforms have the incentive to worsen the type/quality of data released to businesses operating on the platform using the AIAs?

Setting: Sponsored Search Auctions for Digital Ads

Setting: auctions where digital ad space is sold, in particular **Generalized Second-Price (GSP) auctions** used for sponsored search

- Sponsored search auctions: 40% of digital ad revenues
- One dominant platform as a seller (Google), and a few competitors (Bing, Yandex, Seznam, Amazon, etc.) all using auctions
- Buyers increasingly use algorithms (often AIAs) to bid
- Regulation is about to have direct impacts on digital ad services:

Setting: Sponsored Search Auctions for Digital Ads

Setting: auctions where digital ad space is sold, in particular **Generalized Second-Price (GSP) auctions** used for sponsored search

- Sponsored search auctions: 40% of digital ad revenues
- One dominant platform as a seller (Google), and a few competitors (Bing, Yandex, Seznam, Amazon, etc.) all using auctions
- Buyers increasingly use algorithms (often AIAs) to bid
- Regulation is about to have direct impacts on digital ad services:

Digital Markets Act (DMA) 'core platform services' include online ad services (Digital Services Act (DSA) also targets them):

- more restrictions on targeted/micro-targeted ads (DMA Article 6(aa))
- more transparency toward advertisers (DMA Art. 6(g)) and final consumers (DSA Art. 24)

Method and Findings

Method: **controlled experiments via computer simulations**, as in most of the literature (Calvano et al. (2020), Asker et al. (2021), Banchio and Skrzypacz (2022))

Method and Findings

Method: **controlled experiments via computer simulations**, as in most of the literature (Calvano et al. (2020), Asker et al. (2021), Banchio and Skrzypacz (2022))

Findings: the platform can have an **incentive to obfuscate the data** in order to increase revenues

Specifically, we consider the incentive to go

- from a **full information** feedback: all bids revealed
- to **no information** feedback: no competitors' bids revealed

Different information regimes imply the use of different AIAs, so that in our benchmark case platform revenues increase substantially (by more than 20%)

Method and Findings

Method: **controlled experiments via computer simulations**, as in most of the literature (Calvano et al. (2020), Asker et al. (2021), Banchio and Skrzypacz (2022))

Findings: the platform can have an **incentive to obfuscate the data** in order to increase revenues

Specifically, we consider the incentive to go

- from a **full information** feedback: all bids revealed
- to **no information** feedback: no competitors' bids revealed

Different information regimes imply the use of different AIAs, so that in our benchmark case platform revenues increase substantially (by more than 20%)

The baseline result is robust to extensions about both AIAs and auction mechanism

GSP Auction Game

Method: GSP auction game

Series of simulated experiments in which bidders **interact repeatedly** in a **GSP auction for one keyword**

Rules of the GSP auction:

- Bidder i submits a bid $b_i \in \mathbb{R}_+$
- The s-th highest bid (b^s) obtains slot s among several available slots for sale, and pays a price per click equal to b^{s+1}

Method: GSP auction game

Series of simulated experiments in which bidders **interact repeatedly** in a **GSP auction for one keyword**

Rules of the GSP auction:

- Bidder i submits a bid $b_i \in \mathbb{R}_+$
- The s-th highest bid (b^s) obtains slot s among several available slots for sale, and pays a price per click equal to b^{s+1}

Baseline Experiment (3 asymmetric bidders):

- An auction with two slots and three bidders $i \in \{1, 2, 3\}$
- Valuations (per click): $v_1 = 3$, $v_2 = 2$, $v_3 = 1$
- Click-through-rates (CTRs): $x^1 = 5$, $x^2 = 2$
- Feasible bids $\mathcal{B}$ on the interval $[B_{\min,} B_{\max}] = [0.2, 3]$
- $k = 15$ possible bids

Method: GSP auction game

Series of simulated experiments in which bidders **interact repeatedly** in a **GSP auction for one keyword**

Rules of the GSP auction:

- Bidder i submits a bid $b_i \in \mathbb{R}_+$
- The s-th highest bid (b^s) obtains slot s among several available slots for sale, and pays a price per click equal to b^{s+1}

Baseline Experiment (3 asymmetric bidders):

- An auction with two slots and three bidders $i \in \{1, 2, 3\}$
- Valuations (per click): $v_1 = 3$, $v_2 = 2$, $v_3 = 1$
- Click-through-rates (CTRs): $x^1 = 5$, $x^2 = 2$
- Feasible bids $\mathcal{B}$ on the interval $[B_{\min,} B_{\max}] = [0.2, 3]$
- $k = 15$ possible bids

Competitive benchmark equilibrium (EOS) corresponds to:

- bids are equal to $b_1 > b_2$, $b_2 = 1.6$, $b_3 = 1$
- and the platform revenue is $R = 10$

AI Algorithms

Design Features of the AIAs' Auction Experiments: Main Ideas

- Bidders **interact repeatedly** in a **GSP auction for one keyword**
- Each of the bidders uses its own Q-learning algorithm: AIAs learn to bid by **trial and error** in order to maximize the expected present value of the reward stream
- The knowledge of each algorithm is represented by the **Q-matrix**: $Q_t^i(s, b)$, the matrix of **expected rewards** from each possible bid $b \in \mathcal{B}$ in each possible state of the game $s \in \mathcal{S}$ in each period t

Design Features of the AIAs' Auction Experiments: Main Ideas

- Bidders **interact repeatedly** in a **GSP auction for one keyword**
- Each of the bidders uses its own Q-learning algorithm: AIAs learn to bid by **trial and error** in order to maximize the expected present value of the reward stream
- The knowledge of each algorithm is represented by the **Q-matrix**: $Q_t^i(s, b)$, the matrix of **expected rewards** from each possible bid $b \in \mathcal{B}$ in each possible state of the game $s \in \mathcal{S}$ in each period t
- **Updating rule**: after choosing bid b_t^i in state s_t, the algorithm observes r_t^i as well as s_{t+1}, and updates the Q-matrix
  - **Asynchronous** (only for submitted bid) vs
  - **Synchronous** (for all bids by calculating counterfactuals)

Design Features of the AIAs' Auction Experiments: Main Ideas

- Bidders **interact repeatedly** in a **GSP auction for one keyword**
- Each of the bidders uses its own Q-learning algorithm: AIAs learn to bid by **trial and error** in order to maximize the expected present value of the reward stream
- The knowledge of each algorithm is represented by the **Q-matrix**: $Q^i_t(s, b)$, the matrix of **expected rewards** from each possible bid $b \in \mathcal{B}$ in each possible state of the game $s \in \mathcal{S}$ in each period t
- **Updating rule**: after choosing bid b^i_t in state s_t, the algorithm observes r^i_t as well as s_{t+1}, and updates the Q-matrix
  - **Asynchronous** (only for submitted bid) vs
  - **Synchronous** (for all bids by calculating counterfactuals)
- **Data usage**:
  - **Stateful** $s_t = (b^i_{t-1}, b^{-i}_{t-1})$ vs **Stateless** $s_t = \emptyset$

Data Policy: Obfuscation by the Platform ▸ examples

Both the update rule and the definition of the state (i.e., data usage) are not a free choice by the bidder but are a consequence of the platform data policy

Three *information assumptions* describing what the platform reveals about bids:

a. **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 ⇒ **Stateful Synchronous** algorithms

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

c. **No information**: The only information that the bidder observes is the reward she received after submitting a particular bid ⇒ **Stateless Asynchronous** algorithms

Results

Obfuscation Increases the Platform's Revenues

Table: Limit Bids, Rewards, and Auctioneer Revenues

	Bids	Individual Rewards	Revenue
Full Info: Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2
			[7.03, 7.37]
No Info: Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76
			[8.39, 9.13]

The decision not to reveal competitor bids increases the platform's average revenues by 22% from 7.2 to 8.76. Evolution of Revenues, Bids, Rewards

Obfuscation Increases the Platform's Revenues

Table: Limit Bids, Rewards, and Auctioneer Revenues

	Bids	Individual Rewards	Revenue
Full Info: Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2
			[7.03, 7.37]
No Info: Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76
			[8.39, 9.13]

The decision not to reveal competitor bids increases the platform's average revenues by 22% from 7.2 to 8.76. ▸ Evolution of Revenues, Bids, Rewards

- Platform revenue in competitive benchmark of the one-shot game is $R = 10$

Comparison with Other Experimental Designs

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

	Bids	Individual Rewards	Revenue
Full Info: Stateful Synchronous	$(2.03, 1.2, 0.6)$	$(9.0, 2.8, 0.0)$	7.2 [7.03, 7.37]
No Info: Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76 [8.39, 9.13]
Partial Info: 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]

Possible Driver: Implementation of Reward-punishment Schemes

▸ delta

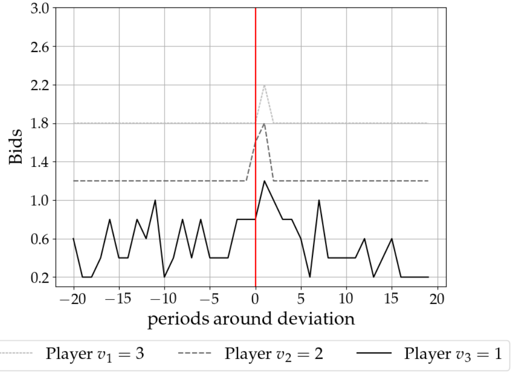

Figure: Evolution of Bids in a Single Run of the Stateful Synchronous Algorithm when Player $v_2 = 2$ Deviates

- 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

Extensions

All of the extensions below lead to differences in the magnitude of the revenue increase via obfuscation, but to the same qualitative outcomes:

- Variation to the GSP game: Milgrom and Mollner (2014) go
- Alternative auction format: VCG mechanism go
- Alternative AIAs: conservative & greedy AIAs via argmax choice; asymmetric grids go

Conclusions

Results from this paper:

1. Data **obfuscation by the platform can improve its revenues and hurt advertisers using AIAs**
2. Algorithmic bidding sustains low bids under the GSP relative to the competitive benchmark

Broader research agenda:

- Competition between **differentially informed AIAs**: platform sponsored AIAs bidding services accessing more/better data
- **Other forms of data obfuscation**: keyword (broad match) and click attribution
- Assessment of DMA-DSA provisions on the interaction between data obfuscation and AIAs bidding

Appendix

Data Policy: Obfuscation Strategies by the Platform ▸ back

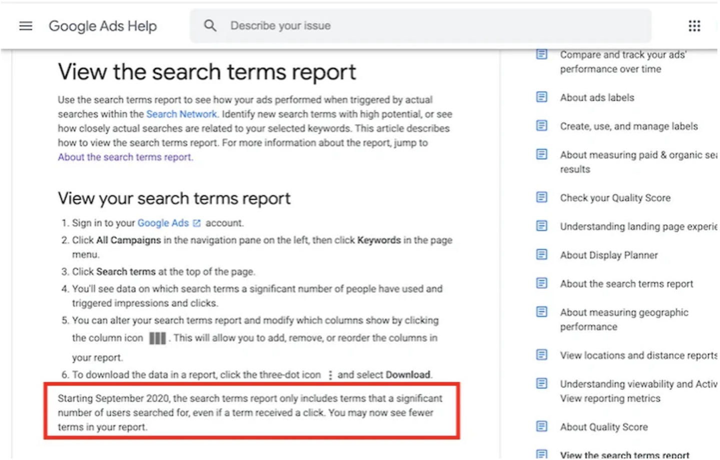

Google Ads Help

Describe your issue

View the search terms report

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

View your search terms report

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

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

Figure: Search Term Report

Data Policy: Obfuscation Strategies by the Platform

▸ back

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

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

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

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

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

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

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

Unsurprisingly, the move has angered advertisers.

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

Figure: Impacts of the Search Term Report Change

Data Policy: Obfuscation Strategies by the Platform

back

About changes to phrase match and broad match modifier

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

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

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

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

Figure: Broad Match Modifiers

Data Policy: Obfuscation Strategies by the Platform ▸ back

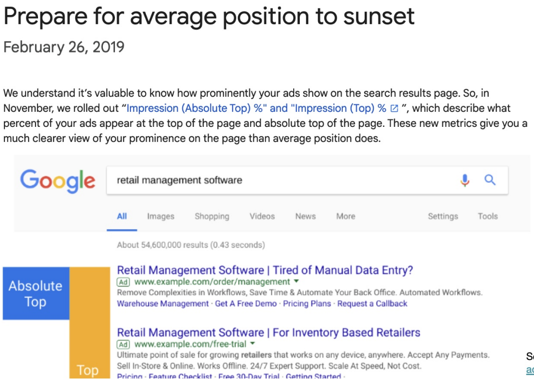

Prepare for average position to sunset

February 26, 2019

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

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

Figure: Position Information

Data Policy: Obfuscation Strategies by the Platform

back

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

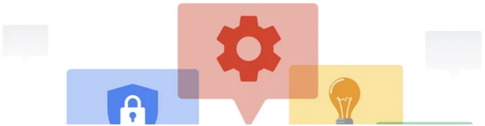

Screenshot taken from the URL https://blog.google/products/ads-commerce/data-driven-attribution-new-default/ on Feb. 20, 2023

Figure: Attribution

Data Policy: Obfuscation Strategies by the Platform

▸ back

Analytics Help

Describe your issue

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

Figure: Details on the Attribution Model

Evolution of Auctioneer's Revenues ▸ back

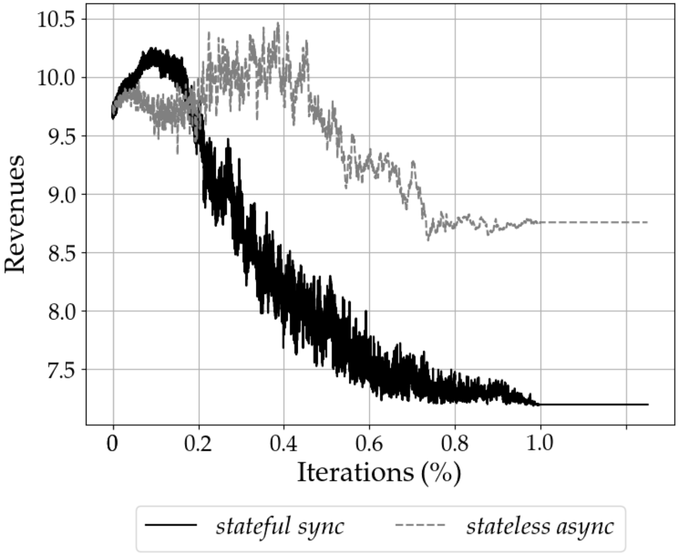

Evolution of Bids and Bidder Rewards

back

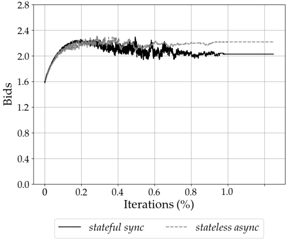

(a) Player $v_1 = 3$ Bids

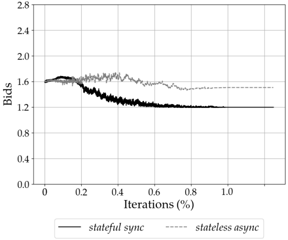

(b) Player $v_2 = 2$ Bids

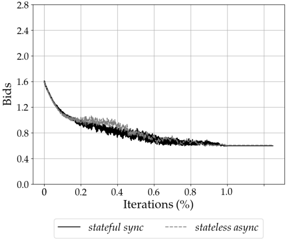

(c) Player $v_3 = 1$ Bids

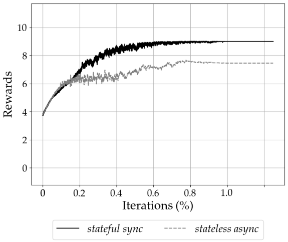

(a) Player $v_1 = 3$ Rewards

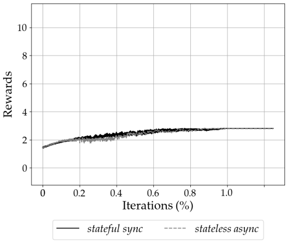

(b) Player $v_2 = 2$ Rewards

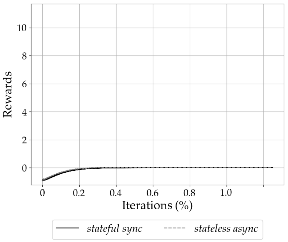

(c) Player $v_3 = 1$ Rewards

Possible Driver: Implementation of Reward-punishment Schemes

back

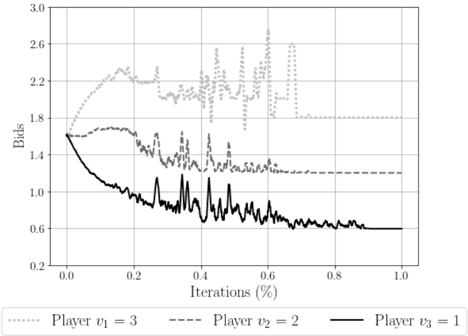

(a) Bids when $\delta = 0.95$

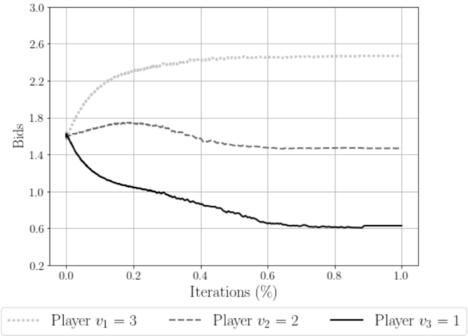

(b) Bids when $\delta = 0$

Figure: Evolution of Bids in a Single Run

Alternative AIAs: Conservative and Greedy via Argmax Selection

back

- If in the exploration process, AIAs are conservative and choose the smallest b among those leading to the highest value of the Q-matrix in a given s, obfuscation leads to an increase in the platform's average revenues by 23% from 6.4 to 7.86
- In the case of greedy algorithms that choose the highest bid, the decision not to reveal competitor bids increases the platform's average revenues by 25% from 8 to 10

	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]

Alternative AIAs: Asymmetric Bid Grids

back

Simulation Results in Case when Player $v_2 = 2$ 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]

One Extra Example: Milgrom and Mollner (2014) back

- $v_1 = 15$, $v_2 = 10$, $v_3 = 5$
- $x^1 = 100$, $x^2 = 3$, $x^3 = 1$
- The set of feasible bids $\mathcal{B}$ on the interval $[B_{\min,} B_{\max}] = [1, 15]$
- $k = 15$ possible bids so that the step between the bids is 1
- The EOS equilibrium in this case is given by $b_1 > b_2$, $b_2 = 9.8$, $b_3 = 3.3(3)$, and leads to auctioneer revenue $R = 990$

	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]

82% increase in auctioneer revenue

Alternative Auction Format: GSP vs VCG ▸ back

Table: 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 [7.03, 7.37]	$(2.53, 1.21, 0.6)$	$(10.18, 2.8, 0.0)$	6.02 [5.66, 6.37]
Stateless Asynchronous	$(2.22, 1.51, 0.61)$	$(7.46, 2.78, 0.0)$	8.76 [8.39, 9.13]	$(2.79, 1.94, 0.62)$	$(7.94, 2.76, 0.0)$	8.3 [7.81, 8.79]

- 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