

Final Project: Ad Exchange Simulation

1 Introduction

This final project option is derived from the [TAC AdX Game](#), ultra-simplified versions of which we explored in Labs 8 and 9.

2 Game Description

There are several differences between this version of the game and the earlier versions from the labs, which we explain in-depth in this handout. In short, the game will be extended to 10 days, campaigns will vary in length, so that one agent may seek to fulfill multiple campaigns simultaneously; and, importantly, some campaigns will be allocated endogenously via auctions (and others, exogenously, according to a known distribution).

2.1 Advertising Campaigns and Impression Opportunities

Your agent's main task in this game is to bid on (and hopefully win) **advertising campaigns**, based on which it will again bid on (and hopefully win) **impression opportunities**, namely opportunities to exhibit ads to Internet users as they browse the web.

Each day of the simulation, a random number of Internet users browse the web. These users hail from multiple market segments, of which there are a total of 26, corresponding to combinations of {Male, Female} x {HighIncome, LowIncome} x {Old, Young}. A market segment might target only one of these attributes (for example, only Female) or two (Female_Young) or three (Female_Old_HighIncome). Users' market segments drawn from the distribution described in Appendix A.

For each user, a **second-price sealed-bid auction** is run to determine which agent to allocate that user's advertising space to, and at what price. Ties are broken randomly, so if there are two winning bidders in a market segment, each will be allocated (about) half the users in that segment at the price they bid.

Ad networks are motivated to participate in these auctions by their desire to fulfill advertising campaigns. A campaign is a contract of the form, "The ad network will show some number of ads to users in some demographic. In return for these impression opportunities, the advertiser agrees to pay the ad network said budget." More specifically, each campaign is characterized by:

- A **market segment**: a demographic(s) to be targeted.
- A **reach**: the number of ads to be shown to users in the relevant market segment.
- A **budget**: the amount of money the advertiser will pay if this contract is fulfilled.
- A **start day** and an **end day**: the time during which the campaign is active. This range is inclusive; you will have the opportunity to bid on a campaign on its start day, its end day, and every day in-between.

Here is an example of a campaign:

```
[Segment = Female_Old, Reach = 500, Budget = $40.0, Start_Day = 4, End_Day = 6]
```

To fulfill this campaign, your agent must show at least 500 advertisements to older women between days 4 and 6. If successful, it will earn \$40. To show an advertisement to a user, you must win the auction for that user. (Yes, we, as Internet users, are regularly auctioned off!) But note that winning an auction for a user who does not match a campaign's market segment does not count toward fulfilling that campaign.

2.2 Ad Auctions: Bids and Spending Limits

Unlike the sealed-bid auctions we have studied in class, which are one-shot auctions, the auctions in this game are repeated, as users arrive repeatedly. However, agents can only bid in these auctions once!—at the start of each day, before simulating them begins. Consequently, agents must reason in advance about how events might unfold over the course of the day, and perhaps make contingency plans. The AdX game provides a mechanism for making a contingency plan in the form of spending limits. These limits are upper bounds on the amount an agent is willing to spend in either a specific market segment, or in total on a campaign, meaning overall, across all market segment associated with that campaign.

If your agent wins a campaign whose market segment is very specific (e.g., `Female_Old_HighIncome`), then it won't really have a choice about which users to bid on; it has to bid for users in precisely that market segment, or it cannot earn a positive profit. However, if its market segment is less specific (e.g., `Female`), it can bid different amounts in the `Female_Old` and `Female_Young` markets, for example, based on how much competition it thinks there will be in each. Keep in mind, though, that the order in which users arrive is random. So if it bids more on `Female_Old` than `Female_Young`, but then if all `Female_Old` users arrive before any `Female_Young`, it may end up spending its entire budget for that campaign on `Female_Old` users. For this reason, when bidding on a market segment, your agent might want to specify a spending limit in each market segment. An agent can also specify a campaign spending limit to ensure that it does not spend more than some preset total across *all* market segments associated with a campaign.

In sum, the key decisions an agent must make in the ad auctions are what bids to place on what market segments, and what spending limits should accompany those bids.

2.3 Effective Reach and Quality Score

At the end of each day, after all of the users have browsed the web and all of the ad auctions have been run, the server will calculate the **effective reach** of all the agents for each of their active campaigns. This value captures how much of a campaign has been fulfilled via a sigmoidal function that relates the number of matching impressions won to the campaign's total reach.

Effective reach is used to calculate an agent's immediate profit on a campaign. Given a campaign C with budget B , cost K , and effective reach $\rho(C)$, profit is calculated as $\rho(C)B - K$. Thus, achieving a high effective reach, in addition to keeping costs low, is key to maximizing profits.

The specific function for the effective reach of a campaign C is given by:

$$\rho(C) = \frac{2}{a} \left(\arctan \left(a \left(\frac{x}{R} \right) - b \right) - \arctan(-b) \right),$$

where $a = 4.08577$, $b = 3.08577$, x is the amount of impressions achieved, and R is your campaign's reach. This function has already been implemented and is available for your use in the code (see Section 4).

Figure 1 plots the effective reach function $\rho(C)$ of two campaigns, one with reach $R = 500$ (blue), and the other with reach $R = 1000$ (red). Note that $\rho(0) = 0$, $\rho(R) = 1.0$, and $\lim_{x \rightarrow \infty} \rho(C) = 1.38442$. The plot shows that the value of obtaining the first few impressions on the road to fulfilling a campaign is relatively low, compared to the value of obtaining the middle and final impressions.

Effective reach is also used to calculate **quality score**, which is a measure of an agent's reputation (used in the campaign auctions), and hence impacts its long-term profits. Quality score is initialized to 1, and then updated at the end of each day as a moving average of the average effective reach, say $\bar{\rho}$, of all campaigns that end on that day. More specifically, an agent's quality score Q is iteratively updated as follows: $Q_{\text{after}} = (1 - \alpha)Q_{\text{before}} + \alpha\bar{\rho}$. Here, α controls how quickly the quality score changes with each day's effective reach.

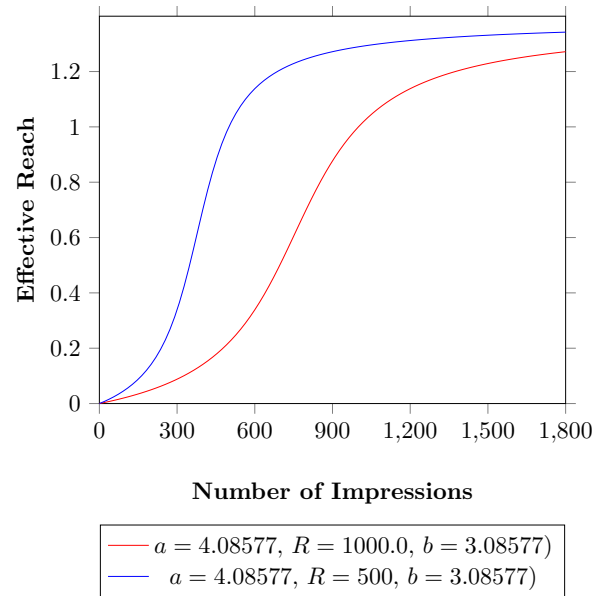


Figure 1: Effective Reach when Reach, R , equals 500 (blue) and 1000 (red).

2.4 Campaign Auctions

At the beginning of the game, your agent will be assigned one campaign at random, whose budget equals its reach. On each subsequent day, with probability $p = \min(1, Q_{\text{agent}})$, your agent will be assigned another random campaign, again with budget equal to reach.

Additionally, each day, multiple (randomly-generated) campaigns will be put up for auction. Campaign auctions are conducted as **second-price reverse auctions**. As opposed to a traditional auction where the good goes to the highest bidder, *in a reverse auction, the contract goes to the lowest bidder*. Reverse auctions are common in procurement. Think of the “bids” as “budgets” for the contract; the lowest bidder is the one who is willing to take on the contract for the lowest budget.

Your agent should submit bids for each campaign it is interested in. These bids, for a campaign with reach R , must fall within the range $[0.1R, R]$.¹ The bids are then divided by the agents’ respective quality scores to arrive at **effective bids**, before they are entered into the auction. For example, for a campaign with reach 100, if two agents bid at the bottom of the acceptable range (i.e., 10), and one’s a quality score is 0.9, and the other’s is 1.1, then their effective bids are $10/0.9 > 10/1.1$, respectively, so the winner—the lower(!) bidder—is the agent with the higher quality score. The budget is then set to $(10/0.9)(1.1)$; in other words, the second-lowest budget scaled by the winning agent’s quality score. If there is only one bidder, say i , in the auction, i wins the campaign with a (maximum) budget of $(R/Q_{\text{low}})Q_i$, where Q_{low} is the average quality score among the three agents with the lowest quality scores on that day, and Q_i is bidder i ’s quality score.

2.5 Scores

At the end of each simulation, the server will compute the profit earned by each agent/ad network. The agent with the highest total profit wins. Because of the randomness inherent in the game, it will be simulated repeatedly, and scores averaged over multiple simulations, to determine an overall winner.

¹The non-zero lower bound on bids is intended to avoid bidding wars, where all agents’ effective bids are zero.

2.6 Putting It All Together

All in all, here is a summary of the AdX game:

Algorithm 1 Run the n -day AdX game

```

Assign each agent a random initial campaign
Set each agent's quality score to 1
for each day  $1 \dots n$  do
  Put  $m$  new campaigns up for auction
  for each agent  $a$  do
    Solicit bids for the new campaigns that are up for auction (get_campaign_bids())
    Solicit bids on market segments for  $a$ 's currently active campaigns (get_ad_bids())
  end for
  Simulate user arrivals and run ad auctions (i.e., allocate users to the winning campaign)
  for each campaign that ends today do
    Compute the relevant agent's effective reach
    Compute that agent's profit (effective reach less cost)
  end for
  Run campaign auctions to allocate the  $m$  new campaigns
  Update the agents' quality scores
end for
The agent with the highest total profit across all  $n$  days wins

```

2.7 Game Specifics

- The game will last for 10 days.
- The quality score adjustment rate α will be 0.5.
- Initial campaigns will be assigned uniformly at random.
- Each day, there will be 5 new campaigns up for auction, also chosen uniformly at random.
 - When the server draws these new campaigns, it removes all campaigns that end after day 10. So, towards the end of the game (especially, on days 9 and 10), there is a (good) chance that there will be fewer than 5 new campaigns.
 - Additionally, each agent will receive a free one-day campaign with a random market segment and a budget equal to its reach with probability $\min(1, Q)$.
- The game will consist of 10 agents, so your agent will have 9 competitors.

The reach and start and end days of campaigns are drawn from distributions (see Appendix A). A campaign's budget (i.e., potential revenue for winning impressions) is set during a campaign auction (see Section 2.4). TAC AdX is a game of incomplete information, as each agent knows (with certainty) the budgets of its own campaigns only, not those of its competitors.

3 Strategy and Considerations

This is a complex game, which requires several strategic considerations beyond what was in play during Labs 8 and 9.

First, your agent can have multiple active campaigns at once, and via its bids in campaign auctions, has some control over which campaigns it is (or at least hopes to be) allocated. If your agent already owns a

campaign for MALE_YOUNG, should it bid low for a campaign for MALE_OLD, so as to not compete too heavily with itself for MALE users? Or should it try to corner the market?

Second, the reverse auction mechanism, and how it determines a campaign's budget, introduces some interesting considerations. How low is your agent willing to bid? If it bids too low, it may end up stuck with a low budget, limiting your agent's profits. This is exacerbated by the fact that the first part of the effective-reach curve is essentially flat—profits don't pick up until about half of a campaign's reach is achieved. On the other hand, if your agent bids too high, it may have trouble winning any campaigns in the first place, and it cannot profit at all without any campaigns.

Third, if your agent wins two campaigns with overlapping market segments, which should it prioritize? Should it place higher bids for the campaign that ends sooner, since it has less time to fulfill that campaign? Or should it pay equal mind to the campaign that ends later, as fulfilling that campaign early would allow it to focus on other overlapping campaigns in the future? Perhaps it can balance its interests by setting precise spending limits based on how much of each campaign its already fulfilled?

If your agent procures only a small percentage of a campaign's reach, then its effective reach for that campaign will be low, which means your agent's profit on that campaign will be low, since the lower part of the effective reach curve is basically flat. But worse still, its quality score will decrease, which means that all else being equal, it must bid lower to win future campaign auctions, so it will have to settle for lower budgets on its future campaigns. In short, your agent must choose wisely when deciding how to bid on campaigns, because it cannot safely ignore any of its active campaigns. Once it wins a campaign, it is stuck with it.

A successful agent strategy for this game will comprise these ideas and many others. We suggest building a theoretical model of the game (consider starting your writeup!) before attempting to implement an agent strategy. This model will help you envision even more of the many important factors that your agent's strategy should try to take into account.

4 Tier 1 Agent

We have provided you with access to our Tier 1 TA bot implementation to test your own agent against. This bot is implemented in the `Tier1NDaysNCampaignsAgent` class, and does the following:

- In all campaign auctions, it bids a random number within the range of valid bids.
- In ad auctions, for each campaign with budget B , reach R , so-far cumulative reach of R_a , and so-far cumulative cost of K_a , it bids $\max(0.1, (B - K_a)/(R - R_a))$ with a spending limit of $\max(1, B - K_a)$.

5 Testing

To test your agent *offline* against other agents, run your `MyNDaysNCampaignsAgent` file directly. The main method instantiates 10 agents, including your own, and runs a local simulation, offline. This simulation consists of 500 iterations of the AdX game.

As with any programming project, you should test your programs extensively. In addition to the usual unit tests, etc., you should play test games against other agents. You can test against 9 copies of your own agent, 9 copies of our Tier 1 TA Agent (see Section 4), or any combination thereof. You can also create your own dummy agents to test your own agent against; to do so, just extend `NDaysNCampaignsAgent`.

6 Submission

Before submitting your code through gradescope, please make sure that

- The submitted agent in `my_ndays_ncampaign_agent.py` is named
- You don't name any of your files `random.py` or `numpy.py` or `<Insert Common Package Name>.py`
- Please let us know if you want to import any new packages not natively provided and we will install them on our end so that your code can import them in your final submission. (E.g., for Lemonade, we installed `tensorflow` for some agents' `tensorflow` code to run correctly.)

A Campaign and User Distributions

Each campaign lasts between 1 and 3 days (chosen uniformly at random), and targets one of 20 possible market segments, a combination of at least two attributes (also chosen uniformly at random). A campaign's reach is given by the average number of users in the selected segment (listed in the tables below), scaled by the campaign's duration in days, times a random reach factor, selected from the set $\{\delta_1, \delta_2, \delta_3\}$, where $0 \leq \delta_i \leq 1$, for all i . The exact values of these factors are tailored to the number of agents in the game. In particular, for the ten-agent games we plan to run, we will use $\delta_1 = 0.3$, $\delta_2 = 0.5$, and $\delta_3 = 0.7$.

Table 1: User Frequencies

Segment	Average Number of Users
Male_Young_LowIncome	1836
Male_Young_HighIncome	517
Male_Old_LowIncome	1795
Male_Old_HighIncome	808
Female_Young_LowIncome	1980
Female_Young_HighIncome	256
Female_Old_LowIncome	2401
Female_Old_HighIncome	407
Total	10000

Table 2: User Frequencies: An Alternative View

	Young	Old	Total
Male	2353	2603	4956
Female	2236	2808	5044
Total	4589	5411	10000

	Low Income	High Income	Total
Male	3631	1325	4956
Female	4381	663	5044
Total	8012	1988	10000

	Young	Old	Total
Low Income	3816	4196	8012
High Income	773	1215	1988
Total	4589	5411	10000

B Game API

B.1 MyNDaysNCampaignsAgent class

Your task is to fill in the following methods of the `MyNDaysNCampaignsAgent` class:

- `get_ad_bids()`: returns a `Set<BidBundle>` representing the agent's bid bundles on the current day. A bid bundle includes a campaign and a spending limit, as well as bids in market segments corresponding to the campaign (and spending limits for those bids as well).
Note that these bids are *daily* bids; they only last for one day. If an agent has a multiple-day campaign, and it places a bid on the first day of its campaign, it must to place another bid on the second day, if it would like to continue bidding. You should feel free to take advantage of this flexibility; it allows your agent's strategy to adjust to the events of yesterday when planning for tomorrow.
- `get_campaign_bids(campaigns_for_auction: Set[Campaign])`: returns a `Dict[Campaign, Float]` representing the agent's bid on each element of `campaigns_for_auction` on which it is bidding.

The server, in playing the role of an ad exchange, calls these two methods every day of the game. Whatever your agent returns is then used by the server when it runs the auctions.

In addition, we have provided the following utility method, which you may fill in, should you find it useful:

- `on_new_game()`: called at the beginning of each simulation of the AdX Game. If you are maintaining any per-game data, this is where you would want to initialize or reset it. For example, if you are using a pre-trained agent, this is where you would read in your saved data. Or, if you are keeping more fine-grained data than that provided by the inherited methods (for example, a map from market segment to the amount of campaigns you have covering it), you would clear it here so that it is empty going into the next game.

Finally, below is a list of methods inherited by `MyNDaysNCampaignsAgent` which you may also find useful:

- `get_current_day() -> int`: gets the current day within the game, a number between 1 and 10.
- `get_active_campaigns() -> Set[Campaign]`: gets the agent's active campaigns. This set should correspond to the campaigns the agent bids on in `get_ad_bids()` (assuming it doesn't ignore any). Only the campaigns owned by your agent will be in this set.
- `get_quality_score() -> float`: gets the agent's current quality score.
- `get_cumulative_reach(campaign: Campaign) -> int`: gets the number of impressions that the agent has won for campaign `campaign so far`. This information is particularly useful for multi-day campaigns; the agent can use it to set goals and spending limits as a campaign proceeds.
- `get_cumulative_cost(campaign: Campaign) -> float`: gets the amount the agent has spent on impressions in the ad auctions for campaign `c so far`. This information is particularly useful for multi-day campaigns; the agent can use it to set goals and spending limits as a campaign proceeds.
- `get_cumulative_profit() -> float`: gets the agent's cumulative profit from the start of the game up to the current day.
- `effective_reach(x: int, R: int) -> float`: the effective reach function.
- `is_valid_campaign_bid(campaign: Campaign, bid: float) -> bool`: checks whether `bid` is in the range $[0.1 * \text{campaign.reach}, \text{campaign.reach}]$. We recommend using this function to verify that your agent's bids are valid before submitting them, as invalid bids are rejected by the server.
- `clip_campaign_bid(campaign: Campaign, bid: float) -> float`: returns `bid`, clipped into the range of valid bids on campaign `campaign`. In other words, returns $\max(0.1 * \text{campaign.reach}, \min(\text{campaign.reach}, \text{bid}))$.
- `effective_campaign_bid(bid: float) -> float`: returns the *effective bid* that would be derived from bidding `bid` in a campaign auction. In other words, returns $\text{bid} / \text{self.get_quality_score}()$.

B.2 The MarketSegment Class

`MarketSegment` is implemented as an extension of `frozenset` with class method `MarketSegment.all_segments()` which returns a list of all the market segments. To iterate over all segments, you can do something like:

```
for segment in MarketSegment.all_segments():
```

To find if segment `m1` is a subset of segment `m2` you can do `m1.issubset(m2)` which returns a `boolean` indicating whether `m1` is a subset of `m2`, so that users in market segment `m1` are also in market segment `m2`. (N.B. market segments are subsets of themselves.)

B.3 Campaign objects

In both the ad and campaign auctions, you will be working with `Campaign` objects. They have the following properties, which for campaign `c` can be accessed as follows:

- `c.uid`
- `c.start_day`
- `c.end_day`
- `c.target_segment`
- `c.reach`
- `c.budget`

B.4 BidBundle and Bid objects

To submit ad bids, you should construct an `Set[BidBundle]`.

Each `BidBundle` represents an agent's bids on a single campaign, and an overall campaign spending limit. `BidBundle` objects are constructed as follows:

```
bundle = BidBundle(
    campaignId: int,
    limit: float,
    bid_entries Set[Bid]
)
```

The `limit` field sets a spending limit campaign-wide; that is, it represents a daily spending limit for the campaign. So even if an agent's market-segment-specific limits (described below) total above this limit, its spending on this campaign will still be capped at this limit.

The `Bid` field allows an agent to set daily bids and spending limits for specific market segments within a campaign. `Bid` objects are constructed as follows:

```
bid = new Bid(
    bidder: NDaysNCampaignsAgent,
    auction_item: Set[str],
    bid_per_item: float,
    bid_limit: float
)
```


So, a `bid_entries` set should contain a `Bid` object for each market segment on which the agent is bidding. Likewise, a `bundle` set should contain one of these objects for each campaign on which the agent is bidding.

C Loading from Saved Data

If you design a strategy that requires training, you can submit a pretrained agent, by which we mean an agent that reads a precomputed data file. To use a pretrained agent that reads its input from a file, you should save the file locally using the provided `path_from_local_root` function in `path_utils.py` (if you don't see this update your stencil with git pull). Here is an example of how to use it in your `update` function, for example, if you save and read from `text.txt`:

```
1 def update(self):
2     # To save a file locally to test.txt, for example, please use the path_from_local_root
3     # method so this filepath is preserved later on when run in the actual competition
4     self.filename = path_from_local_root('test.txt')
5     weight = random.random()
6
7     # Writing to the file
8     with open(self.filename, 'w') as file:
9         [Insert your file writing code here]
10
11    # Reading from the file
12    with open(self.filename, 'r') as file:
13        [Insert your file reading code here]
```

Furthermore, if you specify the `local_save_path` in your `LSVMarena`, every single arena run will be saved as a `.json.gz` file that can be read later to train an agent. A short code snippet for how you might be able to interface with these files is provided in the `process_saved_game` and `process_saved_dir` functions.