

Competitive Safety Strategies in Position Auctions

Danny Kuminov and Moshe Tennenholtz

¹ dannykv@tx.technion.ac.il

² moshet@ie.technion.ac.il

Technion – Israel Institute of Technology,
Haifa 32000, Israel

Abstract. We attempt to address the challenge of suggesting a useful bidding strategy to an agent in the an ad auction setting. We explore the possibility of using competitive safety strategies in that context; a C -competitive strategy *guarantees* a payoff which is no less than a factor of C of the payoff obtained in a best Nash equilibrium. We adopt the model of ad auctions suggested by Varian and provide analysis of competitive safety strategies in that context. We first show that no useful safety competitive strategies exist in a setting with complete information about the agents' valuations. Namely, in a setting with N bidders and exponential click-rate functions the ratio can be as be arbitrarily close to N ; we also show that N is a general upper bound. However, in our main results we show that, surprisingly, useful C -competitive strategies do exist in the incomplete information setting. More specifically, we show that under the assumption that agents' valuations are uniformly distributed, an e -competitive strategy exists for the case of exponential click-rate functions, and a 2-competitive safety strategy exists for linear click-rate functions.

1 Introduction

One of the central challenges of game theory is that of providing a decision maker with an advice about how he should choose his action in a given multi-agent encounter. This challenge, which falls under the so-called prescriptive agenda, has been left without a real answer. For example, the celebrated Nash equilibrium (NE), which is the basis for most game-theoretic analysis, suggests that a multi-agent behavior would be considered “rational” if no decision-maker would prefer to deviate from it, assuming the other decision-makers stick to it. However, while this is a very useful concept from a descriptive point of view, it does not address the question of how *should* a particular agent choose his action in a given game. A NE strategy can only be justified by assuming that the other agents are committed to a specific action profile, which is an unreasonably strong assumption regarding their rationality.

Only very few suggestions have been made in order to address the above challenge. One approach is to suggest to the agent a strategy which will be useful against an opponent taken from a particular class (see e.g. (Powers and Shoham,

2004)). A related idea is to try and learn the opponent model in a repeated interaction in order to optimize behavior against it (Carmel and Markovitch, 1999). Recently, the use of machine learning in order to predict opponent behavior in a game given his behavior in other games, and other agents behaviors, have been shown to lead to significant success (Altman *et al.*, 2006). What is common to the above approaches is that there are no guarantees to our agent, unless we severely restrict the class of opponents he may face. An alternative approach, which is referred to as *competitive safety analysis* has been suggested in (Tennenholtz, 2002) motivated by observation made by Aumann in (Aumann, 1985). This approach deals with guarantees the agent can be provided with, as discussed below.

It is well known that in a purely competitive setting, employing a safety level strategy, one that maximizes the agent's expected utility in the worst case, is the only reasonable mode of behavior. For partially cooperative settings, (Tennenholtz, 2002) justified the use of safety-level strategy by introducing the notion of *C-competitive safety strategy* – a strategy that *guarantees* a payoff which is not less than a factor C of what is obtained in equilibrium. If there exists a C -competitive strategy for small C , then this strategy is a reasonable suggestion for the decision maker. However, the main challenge is whether, for interesting contexts, we do have such competitive safety strategies.

In this work we apply competitive safety analysis to the model of ad auctions, which are mechanisms for assigning online advertisement space to agents according to their (proclaimed) utility from using it. There has been only very limited study of bidding in ad auctions (see e.g. (Borgs *et al.*, 2005; Asdemir, 2005)). Moreover, we are not familiar with any work that deals with the challenging prescriptive problem of how should an agent choose his bids in that setting. The formal model that we use is based on (Varian, 2006), and will be described in the following section. Needless to say, if there exist useful C -competitive safety strategies in the ad auction setting, then they can provide useful means for bidders in such auctions.

The basic model of positions auctions assumes that the bidders' valuations for ad slots are common knowledge. In a more realistic model, each agent knows only his own valuation, while the valuations of all agents are assumed to be taken from some known distribution. We provide an analysis of competitive safety strategies for both the complete and the incomplete information settings. Interestingly, we obtain sharp difference between the usefulness of this approach in these settings. While in the complete information setting, it turns out that no general useful competitive safety strategy exist, they do exist in the (more realistic) incomplete information setting! Namely, we show that in the complete information setting with exponential click-rate functions, assuming N bidders,

the competitive safety ratio can be arbitrarily close to N . We also show that N is a general upper bound. If we assume that the click rates are linear, the ratio can be as bad as $\sum_{t=1}^N \frac{1}{t}$. On the other hand, we show highly positive results in the incomplete information setting. We consider valuations which are taken from the uniform distribution, and two classical types of click-rate functions: the exponential and the linear click-rate functions. We show the existence of an e -competitive strategy for the case of exponential click-rate functions, and a 2-competitive safety strategies for linear click-rate functions.

In section 2 we present and discuss the basic model. In section 3 we provide the analysis of the complete information setting, and in section 4 we provide the analysis of the incomplete information setting.

2 An Ad-Auction Setting

We now provide the model of ad auctions on which we will base our analysis. The formal model that we use is based on (Varian, 2006), with minor changes. The model has originally been presented for the complete information setting, but its adaptation to an incomplete information setting is immediate. The exact assumptions we take when computing safety-level strategies are discussed in sections 3 and 4, when we analyze the complete information and incomplete information settings, respectively.

The ad-auction setting:

- There are N players that compete for S ad slots. It is assumed that $N = S$.
- We denote the *clickthrough rate* (CTR) of a slot by $x_i, i \in \{1 \dots N\}$. The CTR is a publicly known property of a slot, which does not depend on the player who is using it. The slots are numbered in decreasing order of CTR: $\forall i : x_i \geq x_{i+1}$. For ease of presentation, we define $x_i = 0$ for all $i > N$.
- The private value in this model is the utility that each agent derives from a single unit of CTR, which is assumed to be the same regardless of the slot from which it originates. We denote it by $v_i, i \in \{1 \dots N\}$; naturally, $\forall i \in \{1 \dots N\} : v_i > 0$. For ease of presentation, we define $v_i = 0$ for all $i > N$.
- The players' bids are interpreted as the maximal price per unit of CTR they are willing to pay to the CTR provider. We denote them by $\tilde{b}_i, i \in \{1 \dots N\}$; w.l.o.g we assume that $\tilde{b}_i \geq \tilde{b}_{i+1}$; that is, the agents are ordered in decreasing order of bids. Naturally, $\forall i : \tilde{b}_i \geq 0$ by the rules of the auction. For ease of presentation, we define $\tilde{b}_i = 0$ for all $i > N$.
- The mechanism discussed by (Varian, 2006) assigns slots to users according to decreasing order of bids (the highest bidder gets the slot with the highest

CTR, the second highest bidder – the second best slot, etc.). For ease of exposition, we assume that ties are broken according to some predefined ordering of the agents – it can be easily verified that our results hold for any other tie-breaking method as well. The price an agent has to pay per unit of CTR is the bid of the agent immediately below him in this ordering. We denote the price paid by agent i by p_i ; since agents are ordered in decreasing order of bids, $p_i = \tilde{b}_{i+1}$.

- The utility of agent i , which has private value v_i , when the agents' bids are $\tilde{b}_1 \geq \tilde{b}_2 \geq \dots \geq \tilde{b}_i \geq \tilde{b}_{i+1} \geq \dots \geq \tilde{b}_N$ is $(v_i - \tilde{b}_{i+1})x_i$.

A Nash equilibrium is a bidding profile in which each agent prefers his current slot to any alternative slot. Formally:

Definition 1 (A (pure) Nash equilibrium). A Nash equilibrium (NE) is a set of bids $\tilde{b}_1 > \tilde{b}_2 > \dots > \tilde{b}_N$ such that:

1. No agent strictly benefits by decreasing his bid and getting a lesser slot:

$$\forall s, t > s : (v_s - p_s)x_s \geq (v_s - p_t)x_t$$

2. No agent strictly benefits by increasing his bid and getting a better slot:

$$\forall s, t < s : (v_s - p_s)x_s \geq (v_s - p_{t-1})x_t$$

where $p_i = \tilde{b}_{i+1}$.

(Varian, 2006) defined *Symmetric Nash Equilibrium* (SNE) as a bidding profile that satisfies the following:

Definition 2 (SNE). A SNE is a set of bids $\tilde{b}_1 > \tilde{b}_2 > \dots > \tilde{b}_N$ such that:

$$\forall s, t : (v_s - p_s)x_s \geq (v_s - p_t)x_t$$

where $p_i = \tilde{b}_{i+1}$.

Note that the above definitions assume fixed valuations and therefore the game is essentially a *complete information game*.

As shown in (Varian, 2006), SNE has several nice properties:

1. In a SNE, $\forall s : v_s > v_{s+1}$ (i.e. agent i bids higher than agent k iff i 's true type is indeed higher than that of k). This observation is important, since from now on we will assume that agents are indexed in decreasing order of *valuations* and use this property to assert that this order is also the order of their bids in a SNE.

2. If an ordered sequence of bids is a SNE, then it satisfies

$$\tilde{b}_s x_{s-1} \geq v_s(x_{s-1} - x_s) + \tilde{b}_{s+1} x_s$$

3. The latter implies that the bid of agent i in a SNE is bounded as follows:

$$\tilde{b}_i \geq \frac{1}{x_{i-1}} \sum_{t=i}^{N+1} v_t(x_{t-1} - x_t)$$

Note that a bidding profile in which all bids are equal to their respective lower bounds shown above is a SNE; we will use the term *the best SNE* to refer to this bidding profile (the term refers to the players' utility; the "best" SNE actually yields the lowest revenue for the auctioneer). Also, it is a simple observation that in the best SNE, agents never overbid (but there may exist equilibria in which they do so). Finally, it is important to note that the payoff of agent i in the best SNE is $v_i x_i - \sum_{t=i+1}^{N+1} v_t(x_{t-1} - x_t)$, which equals to what his payoff would be in the dominant strategies equilibrium of a VCG auction with the same valuations (his payment is exactly the externality that he imposes on the other players).

Note that it can be helpful to rewrite the expressions (due to (Varian, 2006)) for the bid of agent i (\tilde{b}_i) and the utility of agent i (\tilde{U}_i) in the best SNE as follows:

$$\begin{aligned} \tilde{b}_i &= \frac{1}{x_{i-1}} \sum_{t=i}^{N+1} v_t(x_{t-1} - x_t) = \\ &= \frac{1}{x_{i-1}} [v_i x_{i-1} - v_i x_i + v_{i+1} x_i - v_{i+1} x_{i+1} + \dots \\ &\quad \dots + v_N x_{N-1} - v_N x_N + v_{N+1} x_N - v_{N+1} \cdot 0] = \\ &= \frac{1}{x_{i-1}} \left[v_i x_{i-1} - \sum_{t=i}^N x_t (v_t - v_{t+1}) \right] \\ \tilde{U}_i &= (v_i - \tilde{b}_{i+1}) x_i = v_i x_i - v_{i+1} x_i + \sum_{t=i+1}^N x_t (v_t - v_{t+1}) = \sum_{t=i}^N x_t (v_t - v_{t+1}) \end{aligned}$$

It is important to note that our model slightly differs from that of (Varian, 2006); namely, while the previous work assumes $N > S$, we assume that $N = S$. This models a situation in which the auctioneer has enough ad slots for all the agents, and the only motivation for agents' bidding is the desire to get a higher (better) slot. We think that this is a reasonable model for the online ad auction setting – by the nature of online advertisement, there is no practical limit on

the number of ad slots, and therefore the auctioneer has no real reason to deny agents' requests for slots. For the sake of simplicity, we also assume that there is no reserve price, and therefore the agent with the lowest bid gets the N 'th slot for free. While this is definitely not the case in reality, the results presented here are a good approximation as long as the reserve price is negligible compared to the agents' valuations. Note that although the formulae for the bid and the utility in the *best* equilibrium that are quoted above come from (Varian, 2006), they are true in our model as well (by assuming $v_{N+1} = 0$). This is true since our model can be reduced to that of (Varian, 2006) by adding a fictitious player with fixed valuation 0 – in the best equilibrium, this player always bids truthfully and therefore does not affect slot allocation and expected utility. Note that the reduction does not work for other equilibria, because then the fictitious agent may overbid.

3 Competitive Safety Analysis in the Complete Information Model

Let us consider now the safety level of an agent in the *complete information game* that is induced by the ad-auction presented in the previous section. That is, we assume that the agents' valuations are fixed and are common knowledge, and we want to explore what is the payoff that an agent can guarantee to himself regardless of other agents' behavior. Naturally, if we make no assumptions regarding the rationality of the other agents, they can always force the agent to take the N 'th slot (by bidding higher than the agent's valuation). In this case, it can be easily seen that the payoff loss that the agent suffers relative to his payoff in an equilibrium is unbounded. Even if we limit the other agents not to bid above their valuation³, the agent's payoff in *best* SNE can be $N - \epsilon$ times bigger than his safety level (for any $0 < \epsilon \ll 1$), as can be seen from the following example:

Example 1. Let there be N agents and N slots, let $0 < q < 1$, $x_1 > 0$ be parameters and let the valuations be

$$v_i = \sum_{t=i}^N q^{N-t} = \frac{1 - q^{N-i+1}}{1 - q}$$

and CTR's be

$$x_i = x_1 q^{i-1}$$

That is:

³ This seems a most natural requirement, although there exist Nash equilibria in which agents overbid. It is useful to remember that the rules of ad auction effectively interpret the bid of an agent as the maximal price he is willing to pay per unit of CTR, so by overbidding the agent risks getting negative payoff. Agents do not overbid in the best SNE we compare to.

i	x_i	v_i
1	x_1	$\frac{1-q^N}{1-q}$
2	x_1q	$\frac{1-q^{N-1}}{1-q}$
\vdots	\vdots	\vdots
N-1	x_1q^{N-2}	$q+1$
N	x_1q^{N-1}	1

Let us examine the utility of agent 1. If we assume for a moment that all other agents bid truthfully, a bid by agent 1 that captures slot i results in a payoff of

$$\begin{aligned} U_1 &= (v_1 - v_{i+1})x_i = x_1q^{i-1} \sum_{t=1}^i q^{N-t} = \\ &= x_1q^{i-1} \cdot q^{N-i} \frac{1-q^i}{1-q} = x_1q^{N-1} \frac{1-q^i}{1-q} \end{aligned}$$

This expression is monotone increasing in i ; therefore, agent 1 will prefer to bid 0, take slot N , and get a payoff that equals to $x_1q^{N-1} \frac{1-q^N}{1-q}$. Naturally, the safety level payoff of agent 1 (\hat{U}_1) cannot be higher than this value.

On the other hand, the utility of agent 1 in the best SNE is:

$$\tilde{U}_1 = \sum_{t=1}^N x_t(v_t - v_{t+1}) = \sum_{t=1}^N x_1q^{t-1}q^{N-t} = Nx_1q^{N-1}$$

And the competitive safety ratio is:

$$R(N) = \frac{\tilde{U}_1}{\hat{U}_1} \geq \frac{Nx_1q^{N-1}}{x_1q^{N-1} \frac{1-q^N}{1-q}} = \frac{N(1-q)}{1-q^N}$$

which means that for any N and $0 < \epsilon \ll 1$ we can construct an example in which the competitive safety ratio is at least $N - \epsilon$ by choosing $q \leq \frac{\epsilon}{N}$. \square

In fact, N is an upper bound on the competitive safety ratio, as shown by the following theorem:

Theorem 1. *In the complete information ad auction setting with N slots and N players, the competitive safety ratio (the ratio between an agent's payoff in a best SNE and the payoff guaranteed by a safety level strategy, under the assumption that the agents do not overbid) is at most N .*

Proof. When computing the (pure) safety level of an agent, it is assumed that the other agents know the bid of the agent under consideration (from now on, we will simply refer to him as *the agent*) and all the valuations, and they seek to minimize the utility of the agent. We will use the following notation:

- v_i is the i 'th valuation in the ordered sequence of *all agents'* valuations (including the agent).
- \hat{b}_i is the i 'th bid in the ordered sequence of *all agents'* bids (including the agent).
- v' is the valuation of the agent.
- \hat{b}' is the bid of the agent.
- $v\text{-index}(x) : [0, 1] \rightarrow \{0 \dots N-1\}$ is the number of adversarial agents with valuations that are strictly higher than x (for example, the valuation of the agent under consideration is $v_{v\text{-index}(v')+1}$, the valuation of the adversarial agent immediately below the agent under consideration in the ordering of valuations is $v_{v\text{-index}(v')+2}$, etc.).
- $b\text{-index}(x) : [0, 1] \rightarrow \{0 \dots N-1\}$ is the number of adversarial agents with bids that are strictly higher than x (for example, the bid of the adversarial agent immediately below the agent under consideration in the ordering of bids is $v_{b\text{-index}(b')+2}$).

The utility of the agent is $(v' - \hat{b}_{b\text{-index}(\hat{b}'+2)})x_{b\text{-index}(\hat{b}'+1)}$. Therefore, in order to minimize our agent's payoff, the other agents should choose their bids so that they maximize $\hat{b}_{b\text{-index}(\hat{b}'+2)}$ and $b\text{-index}(\hat{b}') + 1$; those are conflicting goals, since the agents cannot overbid and therefore in order to maximize the price that the agent pays some of the agents with valuation higher than his bid might have to bid lower than him, letting him to get a better slot. It can be easily seen that there is no need to let the agent go up more than one slot and therefore the following is an optimal adversarial strategy:⁴

- All adversarial agents that have valuations smaller than the agent's bid should bid truthfully.
- All adversarial agents that have valuations higher than the agent's bid, *except one*, should bid truthfully.
- One of the players that have higher valuation has to choose whether to bid truthfully or submit the same bid as the agent (which raises the agent's position by one slot but forces him to pay his bid, instead of the valuation of the player below him). Note that we assume here, for ease of exposition, that in the case of a tie the agent under consideration is given the higher slot.⁵

Therefore:

- The agent's utility when all adversarial players bid truthfully is
- $$x_{v\text{-index}(\hat{b}'+1)}(v' - v_{v\text{-index}(\hat{b}'+2)}).$$

⁴ The strategy described here is not unique - we chose the strategy with the most concise description.

⁵ If that is not the case, an approximately equivalent strategy would be to submit a bid that is ϵ -smaller than that of the agent and force him to pay a price that is ϵ -close to his bid.

- The agent’s utility when one of the adversarial players with higher valuation bids the same as he is $x_{v\text{-index}(\hat{b}')} (v' - \hat{b}')$.
- Therefore, the utility of an agent given all valuations and his bid \hat{b} is

$$\min\{x_{v\text{-index}(\hat{b}'+1)}(v' - v_{v\text{-index}(\hat{b}'+2)}), x_{v\text{-index}(\hat{b}')} (v' - \hat{b}')\}$$

Given that the adversarial agents use the strategy described above and the fact that in the complete information setting all the valuations are known to the agent, it can be easily seen that between all the bids that guarantee slot k (i.e. all \hat{b}' so that $v\text{-index}(\hat{b}') = k - 1$), the agent weakly prefers to submit the smallest bid possible. Recall that we assume that ties are decided in favor of the agent⁶, in which case the above observation means that the agent can, without loss of utility, consider only $N - v\text{-index}(v') + 1$ strategies – the valuations $\{v_i : i > v\text{-index}(v') + 1\}$ of the adversarial agents with valuations smaller than his valuation (note that these strategies include 0 as a possible bid which gives the agent the lowest possible slot). Therefore, his safety level payoff is:

$$\begin{aligned} \hat{U} &= \max_{0 \leq \hat{b}'} \min\{x_{v\text{-index}(\hat{b}'+1)}(v' - v_{v\text{-index}(\hat{b}'+2)}), x_{v\text{-index}(\hat{b}')} (v' - \hat{b}')\} = \\ &= \max_{i:v\text{-index}(v')+1 < i \leq N+1} \min\{x_{i-1}(v' - v_i), x_{i-2}(v' - v_i)\} = \\ &= \max_{i:v\text{-index}(v')+1 < i \leq N+1} x_{i-1}(v' - v_i) = \\ &= \max_{i:v\text{-index}(v')+1 \leq i \leq N} x_i(v' - v_{i+1}) \end{aligned}$$

This implies that for any index j such that $v\text{-index}(v') + 1 \leq j \leq N$:

$$\frac{x_j}{\hat{U}} \leq \frac{1}{v' - v_{j+1}}$$

Therefore, the competitive safety ratio is:

$$\frac{\tilde{U}}{\hat{U}} = \frac{\sum_{t=v\text{-index}(v')+1}^N x_t(v_t - v_{t+1})}{\hat{U}} \leq \sum_{t=v\text{-index}(v')+1}^N \frac{v_t - v_{t+1}}{v' - v_{t+1}} \leq N$$

The latter inequality is due to the fact that $\forall t \geq v\text{-index}(v') + 1 : v_t \leq v'$. \square

For the special case when the CTR’s are linear, the worst-case bound is given by the following theorem:

⁶ Otherwise, we would have to consider the set of strategies $\{v_i + \epsilon : i > v\text{-index}(v') + 1\}$ for some small ϵ . It would not affect the nature of the results, but would make the exposition somewhat cumbersome.

Theorem 2. *In the complete information ad auction setting with N slots and N players, when the CTR's are given by $x_i = d(N - i + 1)$ for some $d > 0$, the competitive safety ratio is at most $\sum_{t=1}^N \frac{1}{t} < 1 + \ln N$.*

Proof. The payoff of the agent in the best SNE is:

$$\begin{aligned} \tilde{U} &= (v' - b_{v\text{-index}(v')+2})x_{v\text{-index}(v')+1} = v'x_{v\text{-index}(v')+1} - \\ &- \sum_{t=v\text{-index}(v')+2}^{N+1} v_t(x_{t-1} - x_t) = v'd(N - v\text{-index}(v')) - d \sum_{t=v\text{-index}(v')+2}^{N+1} v_t = \\ &= d \sum_{t=v\text{-index}(v')+2}^{N+1} (v' - v_t) \end{aligned}$$

On the other hand, recall from the proof of theorem 1 that the safety level payoff of the agent is bounded by

$$\begin{aligned} \hat{U} &\leq \max_{i:v\text{-index}(v')+1 \leq i \leq N} x_i(v' - v_{i+1}) = \max_{i:v\text{-index}(v')+1 \leq i \leq N} d(N - i + 1)(v' - v_{i+1}) = \\ &= \max_{i:v\text{-index}(v')+2 \leq i \leq N+1} d(N - i + 2)(v' - v_i) \end{aligned}$$

and therefore for all i such that $v\text{-index}(v') + 2 \leq i \leq N + 1$:

$$(v' - v_i) \leq \frac{\hat{U}}{d(N - i + 2)}$$

Therefore the competitive safety ratio is at most:

$$\begin{aligned} \frac{\tilde{U}}{\hat{U}} &\leq \sum_{t=v\text{-index}(v')+2}^{N+1} \frac{1}{N - t + 2} = \sum_{t=1}^{N-v\text{-index}(v')} \frac{1}{t} = \\ &= 1 + \sum_{t=2}^{N-v\text{-index}(v')} \frac{1}{t} < 1 + \ln N \end{aligned}$$

□

In order to show that $\sum_{t=1}^N \frac{1}{t}$ is a tight bound, let us consider the following example:

Example 2. Let us consider the competitive safety ratio of the agent with valuation $v' = 1$ when the valuations of the other agents are given by $v_i = 1 - \frac{1}{N-i+2}$, and the CTR's are given by $x_i = d(N - i + 1)$ for some $d > 0$. Assuming that all other agents bid truthfully, the payoff to the agent when he captures slot i is:

$$U_i = (v' - b_{i+1})x_i = \frac{1}{N - i + 1} \cdot d(N - i + 1) = d$$

Therefore, his safety level value cannot be greater than d . On the other hand, using the formula from Thm. 2, his payoff in the best SNE is:

$$\tilde{U} = d \sum_{t=2}^{N+1} (v' - v_t) = d \sum_{t=2}^{N+1} \frac{1}{N - t + 2} = d \sum_{t=1}^N \frac{1}{t}$$

And therefore the competitive safety ratio is $\sum_{t=1}^N \frac{1}{t}$. □

4 Competitive Safety Strategies for the Incomplete Information Setting

Now we want to consider the incomplete information setting. Specifically, we want to consider the following decision problem of an agent in this auction:

- The agent under consideration has valuation $v' \in [0, 1]$ that is known to him.
- The agent assumes that other agents' distributions are distributed according to some known distribution.
- The agent is risk-neutral.
- The agent has two possible courses of action:
 1. To select to “play for the best SNE”. This may mean, for example, that the auction is repeated with the same agents (and the same valuations), and the play sequence is assumed to converge to the best SNE. Alternatively, there may exist a central entity (a “mediator”) that offers a course of action that is guaranteed to lead to the best SNE. In any case, the value that the agent assigns to this action is his expected payoff in the best SNE that is induced by the realizations of the players' valuations, where expectation is taken w.r.t the distribution of the other agents' valuations. It is important to note that this is also his expected payoff, given his valuation, in the corresponding Bayes-Nash equilibrium of this auction.⁷

⁷ To see why, consider the following sequence of equalities. The allocation and payments are the same, for any tuple of valuations, in the best SNE and in the corresponding VCG auction (under complete information). However, assuming the incomplete information setup, the VCG

2. To use a “safety level strategy”. This means that the agent selects an action that guarantees him the best expected payoff in the auction, against any *reasonable* action by the other players (where the expectation is taken w.r.t the distribution of the other agents’ valuations). The value that the agent assigns to this action is the guaranteed expected payoff. Specifically, the following model of interaction is assumed:

- All agents are assigned with their private values. Those values are fixed.
- The agent under consideration selects his bid \hat{b}' , based on his valuation only (he does not know the realizations of other agents’ valuations - he only knows their distribution).
- The other agents select their bids based on the bid \hat{b}' and the realizations of *all* valuations (including the agent’s). It is assumed that they can communicate freely. They cannot overbid (i.e. each of them has to submit a bid that is less or equal to his valuation). It is assumed that they select the joint action that minimizes the agent’s utility.

It can be easily seen that in this model of interaction, agents can not gain by using mixed strategies and therefore we can consider, w.l.o.g, only pure strategies.

- Intuitively, all things being equal, action 2 is preferred to action 1, since it does not require elaborate and hard-to-justify assumptions about the rationality of other agents, neither does it require any additional structure on top of the basic auction.
- Therefore, the agent would like to know how much utility he loses by selecting action 2 instead of 1.

Naturally, the answer to the question formulated above depends on the distribution of valuations and the CTR values of the ad slots. We consider uniformly distributed valuations, and two central types of CTRs : exponentially decreasing CTRs, and linearly decreasing CTRs.

4.1 Exponentially Decreasing CTRs

Theorem 3. *In the incomplete information ad auction setting with the following parameters:*

auction is truthful, and therefore for any given valuation, the expected payoff of the agent in the incomplete information setting equals its expected payoff in the corresponding equilibrium of the complete information setting, where the expectation is taken over all possible instantiations of other agents’ valuations. This implies that the expected payoff of our agent in the best SNE, computed according to the realizations of the agents’ valuations, equals its expected payoff in the VCG ad auction with incomplete information. Finally, using the payoff equivalence theorem (Krishna and Perry, 1998) we get that this payoff equals the agent’s payoff, for the given valuation, in the corresponding Bayes-Nash equilibrium of Varian’s mechanism.

- the agents' valuations are distributed independently and uniformly over $[0, 1]$,
- the CTR's are given by $x_k = x_1 q^{k-1}$ for $0 < q < 1$ and $x_1 > 0$,

the ratio between

- the expected payoff in the best SNE that is induced by the valuations' realizations and
- the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid

is at most e .

Proof. Let $v\text{-index}(v')$ be a random variable that represents the number of adversarial agents with valuations that are strictly higher than v' , then the agent's expected utility in the best SNE is:

$$\tilde{U} = \sum_{k=1}^N E(\tilde{U} | v\text{-index}(v') + 1 = k) P(v\text{-index}(v') + 1 = k)$$

$$P(v\text{-index}(v') + 1 = k) = \binom{N-1}{k-1} (1-v')^{k-1} v'^{N-k}$$

$$\begin{aligned} E(\tilde{U} | v\text{-index}(v') + 1 = k) &= E\left(\sum_{t=k}^N x_t (v_t - v_{t+1}) | v\text{-index}(v') + 1 = k\right) = \\ &= x_k (v' - E(v_{k+1} | v\text{-index}(v') + 1 = k)) + \sum_{t=k+1}^N x_t E(v_t - v_{t+1} | v\text{-index}(v') + 1 = k) \end{aligned}$$

Note that given $v\text{-index}(v') + 1 = k$, the valuations $\{v_t : t = k+1, \dots, N\}$ are the order statistics of $N - k$ independent random variables that are distributed uniformly over $[0, v']$ (indexed in decreasing order). Therefore, $v_t \sim v' \cdot \text{Beta}(N - t + 1, t - k)$ and:

$$E(v_t) = v' \frac{N - t + 1}{N - k + 1}$$

$$E(v_t - v_{t+1}) = v' \frac{N - t + 1}{N - k + 1} - v' \frac{N - t}{N - k + 1} = \frac{v'}{N - k + 1}$$

Therefore:

$$E(\tilde{U} | v\text{-index}(v') + 1 = k) = x_k \left(v' - v' \frac{N - k}{N - k + 1} \right) + \sum_{t=k+1}^N x_t \frac{v'}{N - k + 1} =$$

$$= \frac{v'}{N-k+1} \sum_{t=k}^N x_t$$

And therefore the expected utility in best SNE is:

$$\begin{aligned} \tilde{U} &= \sum_{k=1}^N E(U|v\text{-index}(v') + 1 = k)P(i = k) = \\ &= \sum_{k=1}^N \frac{v'}{N-k+1} \left(\sum_{t=k}^N x_t \right) \binom{N-1}{k-1} (1-v')^{k-1} v'^{N-k} = \\ &= \sum_{k=1}^N \frac{1}{N-k+1} \left(\sum_{t=k}^N x_t \right) \binom{N-1}{k-1} (1-v')^{k-1} v'^{N-k+1} = \\ &= \frac{1}{N} \sum_{k=1}^N \left(\sum_{t=k}^N x_t \right) \binom{N}{k-1} (1-v')^{k-1} v'^{N-k+1} \end{aligned}$$

Since

$$\sum_{i=k}^N x_i = \sum_{i=k}^N x_i = x_k \frac{1-q^{N-k+1}}{1-q} = \frac{x_1}{1-q} (q^{k-1} - q^N)$$

the expected utility at equilibrium is:

$$\begin{aligned} \tilde{U} &= \frac{1}{N} \sum_{k=1}^N \left(\sum_{t=k}^N x_t \right) \binom{N}{k-1} (1-v')^{k-1} v'^{N-k+1} = \\ &= \frac{x_1}{N(1-q)} \left[\sum_{k=1}^N \binom{N}{k-1} ((1-v')q)^{k-1} v'^{N-k+1} - \right. \\ &\quad \left. - q^N \sum_{k=1}^N \binom{N}{k-1} (1-v')^{k-1} v'^{N-k+1} \right] \end{aligned}$$

The first sum in the parenthesis is:

$$\begin{aligned} &\sum_{k=1}^N \binom{N}{k-1} ((1-v')q)^{k-1} v'^{N-k+1} = \\ &= \sum_{k=0}^{N-1} \binom{N}{k} ((1-v')q)^k v'^{N-k} = \end{aligned}$$

$$= \sum_{k=0}^N \binom{N}{k} ((1-v')q)^k v'^{N-k} - ((1-v')q)^N = ((1-v')q + v')^N - ((1-v')q)^N$$

Note that the second sum is a particular case of the first sum when $q = 1$, hence:

$$\sum_{k=1}^N \binom{N}{k-1} (1-v')^{k-1} v'^{N-k+1} = 1 - (1-v')^N$$

And therefore:

$$\begin{aligned} \tilde{U} &= \frac{x_1}{N(1-q)} [((1-v')q + v')^N - ((1-v')q)^N - q^N(1 - (1-v')^N)] = \\ &= \frac{x_1}{N(1-q)} [((1-v')q + v')^N - q^N] \end{aligned}$$

On the other hand, in computing the safety level we assume that the other agents know the bid of the agent and all the valuations, and that they seek to minimize the utility of the agent. Recall from section 3 that the utility of the agent in this case, given all valuations and his bid \hat{b}' is

$$\begin{aligned} \min\{x_{v\text{-index}(\hat{b}')+1}(v' - v_{v\text{-index}(\hat{b}')+2}), x_{v\text{-index}(\hat{b}')}(v' - \hat{b}')\} &\geq \\ &\geq x_{v\text{-index}(\hat{b}')+1}(v' - \hat{b}') \end{aligned}$$

where $v\text{-index}(\hat{b}')$ is a random variable that represents the number of adversarial players with valuations strictly higher than the agent's bid). Therefore, the bid that the agent chooses, knowing his valuation and assuming that the other players' valuations are i.i.d with uniform distribution, is the one that maximizes:

$$\hat{U} = \sum_{k=1}^N E(\hat{U} | v\text{-index}(\hat{b}') + 1 = k) P(v\text{-index}(\hat{b}') + 1 = k)$$

$$P(v\text{-index}(\hat{b}') + 1 = k) = \binom{N-1}{k-1} (1-\hat{b}')^{k-1} \hat{b}'^{N-k}$$

$$E(\hat{U} | v\text{-index}(\hat{b}') + 1 = k) \geq x_k(v' - \hat{b}')$$

Therefore:

$$\begin{aligned} \hat{U} &\geq \max_{0 \leq \hat{b}' \leq 1} \sum_{k=1}^N x_k(v' - \hat{b}') \binom{N-1}{k-1} (1-\hat{b}')^{k-1} \hat{b}'^{N-k} = \\ &= \max_{0 \leq \hat{b}' \leq 1} x_1(v' - \hat{b}') \sum_{k=1}^N \binom{N-1}{k-1} ((1-\hat{b}')q)^{k-1} \hat{b}'^{N-k} = \end{aligned}$$

$$\begin{aligned}
&= \max_{0 \leq \hat{b}' \leq 1} x_1(v' - \hat{b}') \sum_{k=0}^{N-1} \binom{N-1}{k} ((1 - \hat{b}')q)^k \hat{b}'^{N-1-k} = \\
&= \max_{0 \leq \hat{b}' \leq 1} x_1(v' - \hat{b}') ((1 - \hat{b}')q + \hat{b}')^{N-1} \\
&= \max_{0 \leq \hat{b}' \leq 1} x_1(v' - \hat{b}') ((1 - q)\hat{b}' + q)^{N-1}
\end{aligned}$$

To find the maximum, we need to compute the derivative and compare it to 0:

$$\begin{aligned}
\frac{d}{d\hat{b}'} \hat{U} &= x_1(-1)((1-q)\hat{b}' + q)^{N-1} + x_1(v' - \hat{b}')(N-1)((1-q)\hat{b}' + q)^{N-2}(1-q) = \\
&= x_1((1-q)\hat{b}' + q)^{N-2} \left[-((1-q)\hat{b}' + q) + (v' - \hat{b}')(N-1)(1-q) \right]
\end{aligned}$$

Since the first part of the expression is always positive, $\frac{d}{d\hat{b}'} \hat{U} = 0$ implies $(v' - \hat{b}')(N-1)(1-q) - (1-q)\hat{b}' - q = 0$ and therefore:

$$\begin{aligned}
v'(N-1)(1-q) - q &= (1-q)\hat{b}' + \hat{b}'(N-1)(1-q) = N(1-q)\hat{b}' \\
\hat{b}' &= \frac{v'(N-1)(1-q) - q}{N(1-q)} = v' \frac{N-1}{N} - \frac{q}{N(1-q)}
\end{aligned}$$

This expression is positive iff $v' \geq \frac{q}{(N-1)(1-q)}$, therefore the maximum is achieved by the following strategy:

$$\hat{b}' = \begin{cases} v' \frac{N-1}{N} - \frac{q}{N(1-q)} & \text{if } v' \geq \frac{q}{(N-1)(1-q)} \\ 0 & \text{otherwise} \end{cases}$$

Now we need to compute the competitive safety ratio in both cases.

Case 1: $v' \geq \frac{q}{(N-1)(1-q)}$

Let us denote $c = ((1 - v')q + v') = ((1 - q)v' + q)$; using this notation, $v' \frac{N-1}{N} - \frac{q}{N(1-q)} = v' - \frac{c}{N(1-q)}$. Therefore, by substituting $\hat{b}' = v' - \frac{c}{N(1-q)}$ we have:

$$\begin{aligned}
\hat{U} &\geq \frac{x_1 c}{N(1-q)} \left((1-q)v' - \frac{c}{N} + q \right)^{N-1} = \\
&= \frac{x_1 c^N}{N(1-q)} \left(1 - \frac{1}{N} \right)^{N-1} \geq \frac{x_1 c^N}{N(1-q)} \cdot \frac{1}{e}
\end{aligned}$$

Recall that:

$$\tilde{U} = \frac{x_1}{N(1-q)} \left[((1-v')q + v')^N - q^N \right] \leq \frac{x_1 c^N}{N(1-q)}$$

And therefore the relative payoff loss from using safety level strategy is bounded in this case by e .

Case 2: $v' < \frac{q}{(N-1)(1-q)}$

In this case, the agent bids $\hat{b}' = 0$, gets the last slot for free and therefore:

$$\hat{U} \geq v' x_1 q^{N-1}$$

On the other hand, recall that

$$\begin{aligned} \tilde{U} &= \frac{x_1}{N(1-q)} [((1-q)v' + q)^N - q^N] = \\ &= \frac{x_1}{N(1-q)} \sum_{i=1}^N \binom{N}{i} ((1-q)v')^i q^{N-i} \end{aligned}$$

And therefore the relative payoff loss from using safety level strategy is bounded in this case by:

$$\begin{aligned} \frac{\tilde{U}}{\hat{U}} &\leq \frac{1}{N(1-q)v'q^{N-1}} \sum_{i=1}^N \binom{N}{i} ((1-q)v')^i q^{N-i} = \\ &= \frac{1}{Nq^{N-1}} \sum_{i=1}^N \binom{N}{i} ((1-q)v')^{i-1} q^{N-i} = \\ &= \frac{1}{Nq^{N-1}} \sum_{i=0}^{N-1} \binom{N}{i+1} ((1-q)v')^i q^{N-i-1} \leq \\ &\leq \frac{1}{q^{N-1}} \sum_{i=0}^{N-1} \binom{N-1}{i} ((1-q)v')^i q^{N-i-1} = \\ &= \frac{1}{q^{N-1}} ((1-q)v' + q)^{N-1} = \left(\frac{(1-q)v' + q}{q} \right)^{N-1} \end{aligned}$$

Recall that in this case $v' < \frac{q}{(N-1)(1-q)}$; replacing v' with this bound in the expression from above:

$$\frac{\tilde{U}}{\hat{U}} < \left(\frac{\frac{q}{N-1} + q}{q} \right)^{N-1} = \left(1 + \frac{1}{N-1} \right)^{N-1} < e$$

Therefore, the payoff loss ratio caused by using a safety level strategy instead of playing for SNE is bounded by e . \square

4.2 Linearly Decreasing CTRs

For the case of uniformly distributed valuations and linearly decreasing CTR's, the answer is given by the following theorem:

Theorem 4. *In the incomplete information ad auction setting with the following parameters:*

- *the agents' valuations are distributed independently and uniformly over $[0, 1]$,*
- *the CTR's are given by $x_k = d(N - k + 1)$ for $d > 0$,*

the ratio between

- *the expected payoff in the best SNE that is induced by the valuations' realizations and*
- *the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid*

is at most 2.

Proof. The utility of agent i in the best SNE is:

$$\tilde{U}_i = (v_i - \tilde{b}_{i+1})x_i = v_i x_i - \sum_{t=i+1}^N v_t (x_{t-1} - x_t)$$

And therefore the expected utility of the agent in the best SNE is:

$$\tilde{U} = \sum_{k=1}^N E(\tilde{U} | v\text{-index}(v') + 1 = k) P(v\text{-index}(v') + 1 = k)$$

$$P(v\text{-index}(v') + 1 = k) = \binom{N-1}{k-1} (1-v')^{k-1} v'^{N-k}$$

$$\begin{aligned} E(\tilde{U} | v\text{-index}(v') + 1 = k) &= E\left(v_k x_k - \sum_{t=k+1}^N v_t (x_{t-1} - x_t) \middle| v\text{-index}(v') + 1 = k\right) = \\ &= v' d(N - k + 1) - d \sum_{t=k+1}^N E(v_t | v\text{-index}(v') + 1 = k) \end{aligned}$$

Note that given $v\text{-index}(v') + 1 = k$, the valuation of each of the $N - k$ players with valuations below v' is uniformly distributed over $[0, v']$. Therefore

the expected value of each of them (without imposing any ordering on them) is $\frac{v'}{2}$ and the expected utility at equilibrium is:

$$E(\tilde{U}|v\text{-index}(v') + 1 = k) = v'd(N - k + 1) - \frac{dv'(N - k)}{2} = dv' \frac{N - k + 2}{2}$$

Hence,

$$\begin{aligned} \tilde{U} &= dv' \sum_{k=1}^N \frac{N - k + 2}{2} \binom{N - 1}{k - 1} (1 - v')^{k-1} v'^{N-k} = \\ &= \frac{dv'}{2} \left[(N + 1) \sum_{k=1}^N \binom{N - 1}{k - 1} (1 - v')^{k-1} v'^{N-k} - \right. \\ &\quad \left. - \sum_{k=1}^N (k - 1) \binom{N - 1}{k - 1} (1 - v')^{k-1} v'^{N-k} \right] \end{aligned}$$

Notice that:

$$\sum_{k=1}^N \binom{N - 1}{k - 1} (1 - v')^{k-1} v'^{N-k} = \sum_{k=0}^{N-1} \binom{N - 1}{k} (1 - v')^k v'^{N-k-1} = 1$$

Similarly,

$$\begin{aligned} \sum_{k=1}^N (k-1) \binom{N - 1}{k - 1} (1 - v')^{k-1} v'^{N-k} &= \sum_{k=2}^N (N-1) \binom{N - 2}{k - 2} (1 - v')^{k-1} v'^{N-k} = \\ &= (N - 1)(1 - v') \sum_{k=0}^{N-2} \binom{N - 2}{k} (1 - v')^k v'^{N-k-2} = (N - 1)(1 - v') \end{aligned}$$

And therefore the expected utility in best SNE is:

$$\begin{aligned} \tilde{U} &= \frac{dv'}{2} ((N + 1) - (N - 1)(1 - v')) = \frac{dv'}{2} (N + 1 - N + Nv' + 1 - v') = \\ &= \frac{dv'}{2} ((N - 1)v' + 2) \end{aligned}$$

On the other hand, recall from the previous theorem that the safety level of the agent is:

$$\hat{U} \geq \max_{0 \leq \hat{v}' \leq 1} \sum_{k=1}^N x_k (v' - \hat{v}') \binom{N - 1}{k - 1} (1 - \hat{v}')^{k-1} \hat{v}'^{N-k} =$$

$$\begin{aligned}
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \sum_{k=1}^N (N - k + 1) \binom{N-1}{k-1} (1 - \hat{b}')^{k-1} \hat{b}'^{N-k} = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \sum_{k=0}^{N-1} (N - k) \binom{N-1}{k} (1 - \hat{b}')^k \hat{b}'^{N-k-1} = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N \sum_{k=0}^{N-1} \binom{N-1}{k} (1 - \hat{b}')^k \hat{b}'^{N-k-1} - \right. \\
&\quad \left. - \sum_{k=1}^{N-1} k \binom{N-1}{k} (1 - \hat{b}')^k \hat{b}'^{N-k-1} \right] = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N - (N-1) \sum_{k=1}^{N-1} \binom{N-2}{k-1} (1 - \hat{b}')^k \hat{b}'^{N-k-1} \right] = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N - (N-1) \sum_{k=0}^{N-2} \binom{N-2}{k} (1 - \hat{b}')^{k+1} \hat{b}'^{N-k-2} \right] = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') (N - (N-1)(1 - \hat{b}')) = \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') ((N-1)\hat{b}' + 1)
\end{aligned}$$

To find the maximum, we need to compute the derivative and compare it to 0:

$$\frac{d}{d\hat{b}'} \hat{U} = -d((N-1)\hat{b}' + 1) + d(v' - \hat{b}')(N-1) = d(N-1)(v' - 2\hat{b}') - d$$

$$\frac{d}{d\hat{b}'} \hat{U} = 0 \Rightarrow (N-1)(v' - 2\hat{b}') - 1 = 0 \Rightarrow \hat{b}' = \frac{v'}{2} - \frac{1}{2(N-1)}$$

Now there are two cases:

Case 1: $\frac{v'}{2} - \frac{1}{2(N-1)} \geq 0 \Rightarrow v' \geq \frac{1}{N-1}$

In this case the the safety level payoff is:

$$\begin{aligned}
\hat{U} &\geq d \left(\frac{v'}{2} + \frac{1}{2(N-1)} \right) \left((N-1) \left(\frac{v'}{2} - \frac{1}{2(N-1)} \right) + 1 \right) = \\
&= d \left(\frac{v'}{2} + \frac{1}{2(N-1)} \right) \left(\frac{v'(N-1)}{2} + \frac{1}{2} \right) = \\
&= \frac{d}{4} \left(v'^2(N-1) + 2v' + \frac{1}{(N-1)} \right)
\end{aligned}$$

And therefore the competitive safety level is:

$$\frac{\tilde{U}}{\hat{U}} \leq 2 \frac{(N-1)v'^2 + 2v'}{v'^2(N-1) + 2v' + \frac{1}{(N-1)}} < 2$$

Case 2: $\frac{v'}{2} - \frac{1}{2(N-1)} < 0 \Rightarrow v' < \frac{1}{N-1}$

In this case, in order to guarantee the maximal value the agent should bid 0, take the lowest slot and get $\hat{U} = v'd$. On the other hand, the expected payoff in the best SNE in this case is (remember that in this case $v' < \frac{1}{N-1}$):

$$\tilde{U} = \frac{dv'}{2} ((N-1)v' + 2) < \frac{3dv'}{2}$$

and therefore the competitive safety ratio is $\frac{\tilde{U}}{\hat{U}} < \frac{3}{2}$. Therefore, the payoff loss ratio caused by using a safety level strategy instead of playing for *SNE* is bounded by 2. \square

4.3 Linearly Decreasing CTRs with quality effect

The actual ad auction mechanism that is used by Google (as described in (Varian, 2006)) uses an additional parameter - the “ad quality” property of an advertiser. It is assumed that the actual amount of clicks that is received by an ad that is allocated to slot i is the product of the slot’s CTR and the quality of the ad. Consequently, Google ranks the ads by the product of a measurement of ad quality and advertiser bid, rather than just the bid alone. The price per click that each advertiser pays is the minimum bid that is necessary to retain his position (given his quality value and the other players’ bids and ad quality values). Since the basic model that we discussed in the previous sections is a special case of this model (derived by assuming that all ads are of equal quality), our negative results regarding the complete information setting apply here as well. To analyze the incomplete information model, we can rely on a simple observation (due to (Varian, 2006)) that this model is equivalent to a variation of the basic model in which the players’ valuations are taken to be their “core” valuations multiplied by the ad quality parameter of the corresponding player. Given this observation, if we assume that both the valuations and the quality of advertisers are distributed independently and uniformly over $[0, 1]$, we have that the equivalent basic model is one in which the valuations are distributed over $[0, 1]$ according to PDF $f(x) = -\ln x$ (Glen *et al.*, 1999). In that case, assuming linearly decreasing CTR’s, the competitive safety ratio is given by the following theorem:

Theorem 5. *In the incomplete information ad auction setting with the following parameters:*

- *the agents' valuations are distributed identically and independently over $[0, 1]$ with PDF $f(x) = -\ln x$*
- *the CTR's are given by $x_k = d(N - k + 1)$ for $d > 0$,*

the ratio between

- *the expected payoff in the best SNE that is induced by the valuations' realizations and*
- *the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid*

is at most

$$1 + \frac{3}{4}(N - 1)$$

Proof. First of all, several properties of the logarithmic distribution. Given random variable Z with PDF $f(x) = -\ln x$:

$$F(a) = P(Z < a) = \int_0^a -\ln x dx = [-x \ln x + x]_0^a = a(1 - \ln a)$$

$$\begin{aligned} E(Z|Z < a) &= \frac{1}{P(Z < a)} \int_0^a x(-\ln x) dx = \\ &= \frac{1}{a(\ln a - 1)} \left[x^2 \left(\frac{\ln x}{2} - \frac{1}{4} \right) \right]_0^a = \frac{a(2 \ln a - 1)}{4(\ln a - 1)} \end{aligned}$$

The utility of agent i in the best SNE is:

$$\tilde{U}_i = (v_i - \tilde{b}_{i+1})x_i = v_i x_i - \sum_{t=i+1}^N v_t (x_{t-1} - x_t)$$

And therefore the expected utility of the agent in the best SNE is:

$$\begin{aligned} \tilde{U} &= \sum_{k=1}^N E(\tilde{U} | v\text{-index}(v') + 1 = k) P(v\text{-index}(v') + 1 = k) \\ P(v\text{-index}(v') + 1 = k) &= \binom{N-1}{k-1} (1 - F(v'))^{k-1} F(v')^{N-k} \\ E(\tilde{U} | v\text{-index}(v') + 1 = k) &= \end{aligned}$$

$$\begin{aligned}
&= E \left(v_k x_k - \sum_{t=k+1}^N v_t (x_{t-1} - x_t) \middle| v\text{-index}(v') + 1 = k \right) = \\
&= v' d(N - k + 1) - d \sum_{t=k+1}^N E(v_t | v\text{-index}(v') + 1 = k)
\end{aligned}$$

Note that given $v\text{-index}(v') + 1 = k$, the expected value of each of the $N - k$ players with valuations below v' (without imposing any ordering on them) is: $E(v | v < v') = \frac{v'(2 \ln v' - 1)}{4(\ln v' - 1)}$ and therefore the expected utility at equilibrium is:

$$\begin{aligned}
E(\tilde{U} | v\text{-index}(v') + 1 = k) &= v' d(N - k + 1) - \frac{d(N - k)v'(2 \ln v' - 1)}{4(\ln v' - 1)} = \\
&= \frac{dv'}{4(\ln v' - 1)} [4(N - k + 1)(\ln v' - 1) - (N - k)(2 \ln v' - 1)] = \\
&= \frac{dv'}{4(\ln v' - 1)} [4N(\ln v' - 1) - (N - 1)(2 \ln v' - 1) - \\
&\quad - (k - 1)(4(\ln v' - 1) - 2 \ln v' + 1)] = \\
&= \frac{dv'}{4(\ln v' - 1)} [2N \ln v' - 3N + 2 \ln v' - 1 - (k - 1)(2 \ln v' - 3)]
\end{aligned}$$

Hence,

$$\begin{aligned}
\tilde{U} &= \frac{dv'}{4(\ln v' - 1)} \left[\sum_{k=1}^N [2N \ln v' - 3N + 2 \ln v' - 1] \right. \\
&\quad \left. \binom{N-1}{k-1} (1 - F(v'))^{k-1} F(v')^{N-k} - \right. \\
&\quad \left. - \sum_{k=1}^N \binom{N-1}{k-1} (1 - F(v'))^{k-1} F(v')^{N-k} (k-1)(2 \ln v' - 3) \right]
\end{aligned}$$

Notice that:

$$\begin{aligned}
&\sum_{k=1}^N \binom{N-1}{k-1} (1 - F(v'))^{k-1} F(v')^{N-k} = \\
&= \sum_{k=0}^{N-1} \binom{N-1}{k} (1 - F(v'))^k F(v')^{N-k-1} = 1
\end{aligned}$$

Similarly,

$$\sum_{k=1}^N (k-1) \binom{N-1}{k-1} (1 - F(v'))^{k-1} F(v')^{N-k} =$$

$$\begin{aligned}
&= \sum_{k=2}^N (N-1) \binom{N-2}{k-2} (1-F(v'))^{k-1} F(v')^{N-k} = \\
&= (N-1)(1-F(v')) \sum_{k=0}^{N-2} \binom{N-2}{k} (1-F(v'))^{k-2} F(v')^{N-k-2} = \\
&= (N-1)(1-v' + v' \ln v')
\end{aligned}$$

And therefore the expected utility in best SNE is:

$$\begin{aligned}
\tilde{U} &= \frac{dv'}{4(\ln v' - 1)} [2N \ln v' - 3N + 2 \ln v' - 1 - \\
&\quad -(2 \ln v' - 3)(N-1)(1-v' + v' \ln v')] = \\
&= \frac{dv'}{4(\ln v' - 1)} [(2 \ln v' - 3)(N-1) + 4 \ln v' - 4 - \\
&\quad -(2 \ln v' - 3)(N-1)(1-v' + v' \ln v')] = \\
&= \frac{dv'}{4(\ln v' - 1)} [4 \ln v' - 4 + v'(3 - 2 \ln v')(N-1)(\ln v' - 1)] = \\
&= dv' \left[1 + \frac{v'}{4}(3 - 2 \ln v')(N-1) \right]
\end{aligned}$$

On the other hand, recall from the previous theorem that the safety level of the agent is:

$$\begin{aligned}
\hat{U} &\geq \max_{0 \leq \hat{b}' \leq 1} \sum_{k=1}^N x_k(v' - \hat{b}') \binom{N-1}{k-1} (1-F(\hat{b}'))^{k-1} F(\hat{b}')^{N-k} = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \sum_{k=1}^N (N-k+1) \binom{N-1}{k-1} (1-F(\hat{b}'))^{k-1} F(\hat{b}')^{N-k} = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \sum_{k=0}^{N-1} (N-k) \binom{N-1}{k} (1-F(\hat{b}'))^k F(\hat{b}')^{N-k-1} = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N \sum_{k=0}^{N-1} \binom{N-1}{k} (1-F(\hat{b}'))^k F(\hat{b}')^{N-k-1} - \right. \\
&\quad \left. - \sum_{k=1}^{N-1} k \binom{N-1}{k} (1-F(\hat{b}'))^k F(\hat{b}')^{N-k-1} \right] =
\end{aligned}$$

$$\begin{aligned}
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N - (N-1) \sum_{k=1}^{N-1} \binom{N-2}{k-1} (1 - F(\hat{b}'))^k F(\hat{b}')^{N-k-1} \right] = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left[N - (N-1) \sum_{k=0}^{N-2} \binom{N-2}{k} (1 - F(\hat{b}'))^{k+1} F(\hat{b}')^{N-k-2} \right] = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left(N - (N-1)(1 - F(\hat{b}')) \right) = \\
&= \max_{0 \leq \hat{b}' \leq 1} d(v' - \hat{b}') \left((N-1)\hat{b}'(1 - \ln \hat{b}') + 1 \right)
\end{aligned}$$

To find the maximum, we need to compute the derivative and compare it to 0:

$$\begin{aligned}
\frac{d}{d\hat{b}'} \hat{U} &= -d \left((N-1)\hat{b}'(1 - \ln \hat{b}') + 1 \right) + d(v' - \hat{b}')(N-1)(-\ln \hat{b}') = \\
&= d(N-1)(-\hat{b}' + \hat{b}' \ln \hat{b}' - v' \ln \hat{b}' + \hat{b}' \ln \hat{b}') - d = \\
&= d(N-1)(2\hat{b}' \ln \hat{b}' + v' \ln \hat{b}' - \hat{b}') - d
\end{aligned}$$

Note that the derivative is always negative and therefore in order to guarantee the maximal value the agent should bid 0, take the lowest slot and get $\hat{U} = v'd$. Therefore the competitive safety ratio is

$$\frac{\tilde{U}}{\hat{U}} \leq 1 + \frac{v'}{4}(3 - 2 \ln v')(N-1)$$

It can be easily verified that as a function of v' , this expression is monotonic increasing over $v' \in [0, 1]$. Therefore the maximal competitive safety ratio is achieved when $v' = 1$, and then

$$\frac{\tilde{U}}{\hat{U}} \leq 1 + \frac{3}{4}(N-1)$$

Therefore, the competitive safety ratio in this case is less or equal to $1 + \frac{3}{4}(N-1)$. \square

4.4 Considering Specific Valuation Values of the Agent

In the first part of this section, we computed bounds on the competitive safety ratio that hold uniformly for all valuations of the agent - now we would like to explore whether those bounds can be improved by limiting the valuation values of the agent. In the terms of the decision problem formulated at the start of this section, showing a good (close to 1) bound for a subset of possible valuations

suggests that if the agent's valuation happens to fall in that subset, he should consider using a safety-level strategy - whereas for other valuations he might do better by using an alternative strategy.

First, we would like to present a simple theorem, which will motivate the rest of this section.

Theorem 6. *In the incomplete information ad auction setting, where the agents' valuations are distributed over $[0, 1]$ with PDF that is well-defined and strictly positive in $[0, 1]$ ⁸, if R is the ratio between*

- *the expected payoff in the best SNE that is induced by the valuations' realizations and*
- *the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid*

, then $\forall N : \lim_{v' \rightarrow 0^+} R = 1$, where v' is the valuation of the agent.

Proof. Let x_N denote the CTR of the last slot and let \hat{U} denote the safety level payoff of the agent, then for any valuation $v > c$ of the agent: $\hat{U} \geq vx_N$ (because the agent can always bid 0 and guarantee at least the last slot). On the other hand, the expected equilibrium payoff is:

$$\begin{aligned} \tilde{U} &= \sum_{k=1}^N E(\tilde{U} | v\text{-index}(v') + 1 = k) P(v\text{-index}(v') + 1 = k) = \\ &= vx_N + \sum_{k=1}^{N-1} E(\tilde{U} | v\text{-index}(v') + 1 = k) P(v\text{-index}(v') + 1 = k) \leq \\ &\leq vx_N + x_1 v P(v\text{-index}(v') + 1 < N) \end{aligned}$$

And therefore:

$$\begin{aligned} \lim_{v' \rightarrow 0^+} R &\leq \lim_{v' \rightarrow 0^+} \frac{vX_N + x_1 v P(v\text{-index}(v') + 1 < N)}{vx_N} = \\ &= 1 + \lim_{v' \rightarrow 0^+} \frac{x_1}{x_N} P(v\text{-index}(v') + 1 < N) = 1 \end{aligned}$$

The latter equality is due to the fact that as the valuation of the agent approaches 0, the probability that all other agents' valuation are greater than his goes to 1. \square

⁸ That is, the distribution assigns zero probability to any singleton subset, and any subset of cardinality \aleph has strictly positive probability.

The above theorem, while important from a motivational point of view, has little practical importance. While it promises that for any desired competitive safety ratio $R > 1$, there exists $c > 0$ such that for all $v' \in [0, c]$ the ratio is at most R , it does not show how to compute this c , neither does it promise that this c is big enough to be practically useful. In order to analyze the behavior of the competitive safety ratio for small values of v' , additional assumptions regarding the distribution of valuations and the CTR values of the ad slots are required; the following observations follow directly from the analysis presented in the first part of this section, for the appropriate models:

Observation 1 *In the incomplete information ad auction setting with the following parameters:*

- the agents' valuations are distributed independently and uniformly over $[0, 1]$,
- the CTR's are given by $x_k = x_1 q^{k-1}$ for $0 < q < 1$ and $x_1 > 0$,

, when the agent's valuation is $v' < \frac{q}{(N-1)(1-q)}$, the ratio between

- the expected payoff in the best SNE that is induced by the valuations' realizations and
- the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid

is given by

$$\left(\frac{(1-q)v' + q}{q} \right)^{N-1} = \left(1 + \frac{1-q}{q} v' \right)^{N-1} \simeq e^{\frac{(N-1)(1-q)}{q} v'}$$

That is, for small valuations of the agent, the competitive safety ratio decreases exponentially as v' approaches 0. While this may sound promising, the relevant range is very small - to illustrate, if we take $q = 0.5$ and $N = 10$, in order to achieve competitive safety ratio of 1.1 the valuation of the agent must be less than 0.01 (approximately).

Observation 2 *In the incomplete information ad auction setting with the following parameters:*

- the agents' valuations are distributed independently and uniformly over $[0, 1]$,
- the CTR's are given by $x_k = d(N - k + 1)$ for $d > 0$,

, when the agent's valuation is $v' < \frac{1}{N-1}$, the ratio between

- the expected payoff in the best SNE that is induced by the valuations' realizations and

- the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid

is at most $\frac{1}{2}((N-1)v' + 2) = 1 + \frac{N-1}{2}v'$.

That is, for small valuations of the agent, the competitive safety ratio decreases linearly as v' approaches 0. To demonstrate, let us take $N = 11$, then in order to achieve competitive safety ratio of 1.1 the valuation of the agent must be less than 0.02.

Observation 3 *In the incomplete information ad auction setting with the following parameters:*

- the agents' valuations are distributed identically and independently over $[0, 1]$ with PDF $f(x) = -\ln x$
- the CTR's are given by $x_k = d(N - k + 1)$ for $d > 0$,

the ratio between

- the expected payoff in the best SNE that is induced by the valuations' realizations and
- the expected payoff guaranteed by the safety level strategy, under the assumption that other agents do not overbid

is at most

$$1 + \frac{v'}{4}(3 - 2 \ln v')(N - 1)$$

To demonstrate, let us take $N = 11$, then in order to achieve competitive safety ratio of 1.1 the valuation of the agent must be less than 0.0027 (approximately). While this value may sound too small to be practical, note that if we assume that the agent's valuation has the same distribution as the opponents' (i.e. its CDF is $P(Z < a) = a(1 - \ln a)$), the prior probability for the agent's valuation to fall in this range is approximately 0.019 - which means that in about 2% of all interactions, the agent can find using the safety level strategy very useful.

5 Conclusions and Future Work

In this work, we have investigated whether useful C-competitive strategies exist in the setting of ad auctions, both in the complete and incomplete information models. We have focused our work on a model in which the slot values are decreasing exponentially or linearly, which we believe to be realistic assumptions.

For these settings, we have shown by examples that in the complete information model there is no hope of achieving constant competitive safety ratio. On the other hand, in the incomplete information model with uniformly distributed valuations a competitive safety ratio of e can be achieved for exponentially decreasing CTRs, and a competitive safety ratio of 2 can be achieved for linearly decreasing CTRs. The intuition behind the difference in results for the complete and incomplete information settings, is that while we show a specific profile of valuations that is arbitrarily bad for the agent, the probability that “bad” profiles actually occur is negligible, and the profiles that do occur with high probability exhibit constant competitive safety ratio.

We see two conceptually different directions for future work:

- Investigating the existence of C -competitive strategies in other models of ad auctions. Those may include non-uniform distributions of valuations, other models of slot values or refinement of auction rules (such as introducing the quality factor parameter (Varian, 2006)).
- Investigating the existence of C -competitive strategies in other interesting sub-classes of games, such as congestion games (Rosenthal, 1973), other non-VCG auctions, etc.

References

- [Altman *et al.*, 2006] A. Altman, A. Boden-Bercovici, and M. Tennenholtz. Learning in one-shot strategic form games. In *Proceedings of ECML-06*, 2006.
- [Asdemir, 2005] K. Asdemir. Bidding Patterns in Search Engine Auctions. Working Paper, University of Alberta School of Business. Available at <http://www.business.ualberta.ca/kasdemir/biddingWARS.PDF>, 2005.
- [Aumann, 1985] R.J. Aumann. On the Non-Transferable Utility Value: A Comment on the Roth-Shaper Examples. *Econometrica*, 53(3):667–677, 1985.
- [Borgs *et al.*, 2005] C. Borgs, J. Chayes, O. Etesami, N. Immorlica, K. Jain, and M. Mahdian. Bid Optimization in Online Advertisement Auctions. manuscript, 2005.
- [Carmel and Markovitch, 1999] David Carmel and Shaul Markovitch. Exploration strategies for model-based learning in multiagent systems. *Autonomous Agents and Multi-agent Systems*, 2(2):141–172, 1999.
- [Glen *et al.*, 1999] A. Glen, L. Leemis, and J. Drew. Computing the distribution of the product of two continuous random variables, 1999.
- [Krishna and Perry, 1998] V. Krishna and M. Perry. Efficient Mechanism Design. Technical report, Pennsylvania State University. Available at <http://econ.la.psu.edu/~vkrishna/papers/vcg20.pdf>, 1998.
- [Powers and Shoham, 2004] R. Powers and Y. Shoham. New criteria and a new algorithm for learning in multi-agent systems. *Proceedings of NIPS 2004*, 2004.
- [Rosenthal, 1973] R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
- [Tennenholtz, 2002] M. Tennenholtz. Competitive Safety Analysis: Robust Decision-Making in Multi-Agent Systems. *Journal of Artificial Intelligence Research*, 17:363–378, 2002.

[Varian, 2006] H. R. Varian. Position auctions. Technical Report, UC Berkeley. Available at <http://www.sims.berkeley.edu/~hal/Papers/2006/>, 2006.