

Optimal Mechanism Design for Agents with DSL Strategies: The Case of Sybil Attacks in Combinatorial Auctions

Yotam Gafni

Technion
Haifa, Israel

yotam.gafni@campus.technion.ac.il

Moshe Tennenholtz

Technion
Haifa, Israel

moshet@ie.technion.ac.il

In robust decision making under uncertainty, a natural choice is to go with safety (aka security) level strategies. However, in many important cases, most notably auctions, there is a large multitude of safety level strategies, thus making the choice unclear. We consider two refined notions:

- a term we call DSL (distinguishable safety level), and is based on the notion of “discrimin” [7], which uses a pairwise comparison of actions while removing trivial equivalencies. This captures the fact that when comparing two actions an agent should not care about payoffs in situations where they lead to identical payoffs.
- The well-known Leximin notion from social choice theory, which we apply for robust decision-making. In particular, the leximin is always DSL but not vice-versa [7].

We study the relations of these notions to other robust notions, and illustrate the results of their use in auctions and other settings. Economic design aims to maximize social welfare when facing self-motivated participants. In online environments, such as the Web, participants’ incentives take a novel form originating from the lack of clear agent identity—the ability to create Sybil attacks, i.e., the ability of each participant to act using multiple identities. It is well-known that Sybil attacks are a major obstacle for welfare-maximization. Our main result proves that when DSL attackers face uncertainty over the auction’s bids, the celebrated VCG mechanism is welfare-maximizing even under Sybil attacks. Altogether, our work shows a successful fundamental synergy between robustness under uncertainty, economic design, and agents’ strategic manipulations in online multi-agent systems.

1 Introduction

Consider an agent who needs to decide on her action in an environment consisting of other agents. In certain cases there is a uniquely defined optimal action for the agent, but in most cases this “agent perspective” is an open challenge. Given the above, both AI and economics care about an adequate modeling of an agent, and its ramifications in a variety of multi agent contexts, for example, on social welfare.

We consider a notion for agent modeling we term DSL (Distinguishable Safety-Level). The notion was previously suggested in the context of constraint-satisfaction problems and fuzzy logic, and was termed “discrimin” [7]. In game theoretic settings, the notion was previously applied [6] as a solution concept for bargaining in Boolean games [11]. To the best of our knowledge, it was not previously considered in the context of auctions, voting, and more generally mechanism design, i.e., when considering the robustness of economic mechanisms’ performance when facing strategic agents.

There are two ways to think of the DSL solution concept, when applied to agent modeling. One is as a solution concept adapted to capture the behavioral phenomenon of the loss aversion cognitive bias in agents, particularly when probabilities over nature states are unknown. The other is as a form of robust strategy choice under uncertainty, that may be required in volatile and unpredictable environments that

do not admit a stable Bayesian description. We show its usefulness in auctions. In the full version of this paper, we also study its behavior in other prominent strategic settings, such as voting. In our main result we consider the celebrated welfare maximizing VCG mechanism in combinatorial auctions setting, where it is known to fail under false name (aka Sybil) attacks. We show that DSL agents lead to optimal social welfare.

1.1 Reasoning under Uncertainty

A classic distinction [15] separates reasoning under risk, where the actors are rational and there is a commonly known distribution about their environment (also known as the stochastic or Bayesian setting), and reasoning under uncertainty, where the general structure of strategies and outcomes is known, but there is no probabilistic information about the environment. Moreover, even assumptions regarding actors' rationality or behavior characteristics may not be guaranteed. For such cases, a robust or worst-case approach seems appropriate, and various notions exist to capture it. Ideally, a dominant strategy solution exists, but this is usually not the case (and indeed it is not the case in all the cases we analyze in this paper). A minimal robust notion is that of a safety level strategy, which uses a max-min approach over all possible outcomes given a strategy choice. However, though it yields interesting results in some cases [2, 20], in many other cases it does not tell us much about what strategy to choose, in particular in auctions settings, where we derive our most interesting results. As we see, this is because in auctions the natural safety level is 0 (which happens when the bidder loses the auction), and any strategy that does not overbid guarantees it. It is thus hard to choose among these strategies without considering a more refined notion. Existing refined notions are the lexicographic max-min (originally defined in [19]) and min-max regret [18]. We overview their comparison to the notion of DSL in Section 3 and Appendix A, respectively.

1.2 VCG, Sybil Attacks, and Welfare

VCG is a well known mechanism which can be applied for combinatorial auctions. VCG has good qualities such as being dominant strategy incentive compatible and achieving optimal social welfare. However, under the possibility of false-name attacks [22], it is no longer truthful. Coming up with other mechanisms does not solve the basic conundrum: In the full information settings, any false-name proof mechanism performs poorly in terms of welfare [12].

A possible avenue to solving the issue is by limiting the discussed valuation classes. However, an example in [13] shows that even when all bidders have sub-modular valuations, VCG is no longer dominant strategy incentive-compatible under false-name attacks. Notably though, even with this example, VCG still arrives at the socially optimal allocation, and in fact as [1] show, this observation is true in general up to a constant with sub-modular (and near sub-modular) bidders. However, in the full version of our paper, we show an example where for the XOS valuation class, which extends the sub-modular class, there is such an attack so that VCG arrives at an arbitrarily sub-optimal allocation. The attack we describe is enabled by the full information settings. Without full information, the attack is risky for the attacker, since it could lead to negative utility, as the attacker overbids her true valuation.

A useful approach, that can lead to better welfare guarantees than dominant strategy mechanism design, is Bayesian mechanism design. Assuming that the bidder distributions are common knowledge, recent work has shown that selling each item separately leads to good constant approximation welfare guarantees for XOS [4] and sub-additive [8] valuations. Though the works do not explicitly consider




Figure 1: Hierarchy for robust decision under uncertainty

false-name attacks, their constructions use the false-name-proof first and second price auctions to auction items separately, and so their results naturally extend to Bayesian false-name mechanism design.

It is important to note, that many of the above positive results for welfare guarantees under false-name attack assume some form of risk-aversion; most importantly, that bidders do not overbid, i.e., they choose only strategies that are individually rational (under any possible nature state). This condition is equivalent to limiting the strategy space only to safety level strategies (as in this case of combinatorial auctions, the safety level is 0). In [10] the authors do not make this assumption, but their positive welfare optimality results are limited as they only consider the homogeneous single-minded case with two items. We thus believe that it is natural to ask: Under our definition of DSL, which is a strong risk-aversion notion (compared, e.g., to the safety level strategy), what welfare guarantees can be obtained? Surprisingly, the answer is optimal, as we show in our main result in Theorem 4.2.

1.3 Our Results

In Section 2 we formally define our solution concept, and apply it to the first-price and discrete first-price auctions. In Section 3 (with the additional discussion of min-max regret in Appendix A) we describe a hierarchy of solution concepts and their relations to the solution concept we introduce (DSL), as summarized in Figure 1.

In Section 4, we present our main result. We discuss VCG as a combinatorial auction under false-name attacks, when bidders may create shell identities to send bids. It is known that VCG is not dominant strategy truthful in these settings, and previous results were limited in establishing good welfare guarantees for combinatorial auctions generally under false-name attacks. We show that when bidders use DSL strategies, VCG achieves optimal welfare even under the threat of false-name attacks.

2 DSL: Definition

When defining DSL strategies, we take the perspective of a single agent i facing uncertainty. The agent has a utility function u_i that determines her utility given the state of the world, which is comprised of her own action a_i , others' actions a_{-i} , and agent i 's type θ_i . Formally, $u_i(a_i, a_{-i} | \theta_i)$. We denote by A_i the set of all agent i 's pure actions, and by $\Delta(A_i)$ the set of all agent i 's mixed actions. An action a_i may be from either of these action sets depending on the context. For mixed strategies, $u_i(a_i, a_{-i} | \theta_i) = \mathbb{E}_{a_{-i}} [u_i(a_i, a_{-i} | \theta_i)]$. We denote by Θ_i the set of all agent i 's types.

Definition 2.1. We say that an action a_i of agent i is *DSL* (given a type θ_i) if for any other action a'_i , over the set of outcomes where agent i 's utility differs between the actions, the minimal utility attained using a_i is at least as good as that attained by a'_i . Formally, let

$$D_{\theta_i}(a_i, a'_i) = \{a_{-i} \text{ s.t. } u_i(a_i, a_{-i} | \theta_i) \neq u_i(a'_i, a_{-i} | \theta_i)\}.$$

Then, an action a_i is pure/mixed DSL if $\forall a'_i \in A_i$,¹

$$\min_{a_{-i} \in D_{\theta_i}(a_i, a'_i)} u_i(a_i, a_{-i} | \theta_i) \geq \min_{a_{-i} \in D_{\theta_i}(a_i, a'_i)} u_i(a'_i, a_{-i} | \theta_i).^2$$

We say that a strategy $s_i : \Theta_i \rightarrow A_i$ is pure DSL if it maps any type θ_i to a corresponding DSL pure action a_i . We say that $s_i : \Theta_i \rightarrow \Delta(A_i)$ is mixed DSL if it maps any type θ_i to a corresponding DSL mixed action a_i .

Notice that in our definition we compare pure strategies only with other pure strategies, i.e., they are DSL with respect to this strategy set. Mixed strategies are DSL w.r.t. all strategies (mixed and pure). We use the term ‘‘nature state’’ to mean the actions a_{-i} , which may result from either uncertainty over others’ types or over their strategic choice: What matters to the agent in the end is what are all of their possible actions. There is seemingly some loss of generality in that we assume that all possible a_{-i} are fixed vectors of actions, and not more generally random variables over actions. But, as we show in the full version of our paper, allowing for the latter loses the usefulness of the DSL notion.

3 Relations to Prominent Game-theoretic Solution Concepts

Note: Missing proofs in this section appear in the full version of our paper. For completeness, we state the connection of DSL to safety level and what we call Multi-Leximin strategies, although these claims are already established in the literature characterizing the notion of discrim (see, e.g., [7]).

3.1 Dominant Strategy

Definition 3.1. A weakly dominant action a_i satisfies that for any other action a'_i :

(1) For any nature state a_{-i} ,

$$u_i(a_i, a_{-i} | \theta_i) \geq u_i(a'_i, a_{-i} | \theta_i),$$

and (2) there is such nature state a_{-i} so that the above inequality is strict.

A weakly dominant strategy is such that maps types to weakly dominant actions.

The following result is natural:

Lemma 3.2. Every weakly dominant strategy is DSL.

3.2 Safety Level Strategy and Individual Rationality

Safety level strategies in non-cooperative games are such strategies that yield a best possible guarantee of utility for a player, without the need to reason about the types or strategies chosen by other players. The example of [2] makes a compelling argument for choosing such strategies: There are games where the Nash Equilibrium does not guarantee more than the safety level. In such cases, choosing the equilibrium strategy runs the unnecessary risk of a lower outcome. [21] extends this insight and shows a class of games where the safety level strategy guarantees a large constant fraction of the Nash equilibrium outcome, without its involved risks.

¹Or, in the mixed case: $\forall a'_i \in \Delta(A_i)$

²We use the term minimum loosely: When taken over infinite sets that do not have a minimum the definition uses the infimum.

Individual rationality is a common requirement in game theory analysis (see, e.g., [17]), that requires either that an agent does not participate in a game where it gains negative utility, or that it does not choose a strategy that may yield negative outcomes. We define:

Definition 3.3. A safety level strategy s_i [2] is a strategy (mixed or pure) of player i such that for any type θ_i it chooses an action a_i so that for any nature state a_{-i} of the other agents, $u_i(a_i, a_{-i} | \theta_i) \geq \max_a \min_{a_{-i}} u_i(a, a_{-i} | \theta_i)$. I.e., the strategy guarantees the safety level $L \stackrel{def}{=} \max_{a'} \min_{a'_{-i}} u_i(a', a'_{-i} | \theta_i)$.

Individual Rationality of a strategy s_i of player i satisfies that for any type θ_i it chooses an action a_i so that for any nature state a_{-i} of the other agents, $u_i(a_i, a_{-i} | \theta_i) \geq 0$. I.e., the strategy guarantees a non-negative utility for the player.

The two notions are quite similar, as individual rationality can be seen as a minimal safety level requirement; in auctions they are in fact equivalent to a third notion of no over-bidding, under some reasonable conditions (the auction does not charge payments from non-winners, and never charges a winner more than her declared value). We claim:

Proposition 3.4. A DSL strategy is a safety level strategy, but not necessarily vice-versa.

Corollary 3.5. When there is a finite amount of safety level strategies, and a finite amount of nature states, a DSL strategy is guaranteed to exist.

The corollary is a result of Lemma 3.8 and Lemma 3.9. We prove both during our discussion of the lexicographic max-min in the next subsection.

3.3 Lexicographic Max-min

A very interesting comparison is with another robust solution notion, the lexicographic max-min (also commonly known as leximin). The leximin is especially prevalent in the fair allocation literature, see, e.g., [16]. We consider two possible ways to define it:

Definition 3.6. Leximin - Let U_{a_i} be the **set** of all possible utility outcomes of the action a_i by agent i , ordered from small to large, and let $U_{a_i}[j]$ be the j element of U_{a_i} in this ordering. An action a_i lexicographically weakly dominates (LD) another action a'_i if $\min U_{a_i} > \min U_{a'_i}$, or $\min U_{a_i} = \min U_{a'_i}$ and $U_{a_i} \setminus \min U_{a_i}$ LDs $U_{a'_i} \setminus \min U_{a'_i}$ (a recursive definition). We call an action that LDs all other actions a leximin. A strategy is leximin if it maps all types to leximin actions.

Multi-Leximin - Let U_{a_i} be the **multiset** of all possible utility outcomes of the action a_i by agent i , ordered from small to large. The rest of the definition follows similarly, where importantly in the recursive definition we remove only one copy of the minimum element at each step.

Note that the (Multi-)leximin notions are only clearly defined when there is a finite amount of nature states a_{-i} , otherwise the recursive definition of LD may not terminate.

We first note that both definitions give stronger notions than safety level strategies.

Lemma 3.7. (Multi-)leximin is a safety level strategy, but not necessarily vice-versa.

Despite some similarity in the definition with DSL, the notion of leximin does not have a special relationship with it: neither implies the other. We demonstrate it using the discrete first-price auction in Example B.7 Appendix B.

The notion of multi-leximin is much more closely related to the DSL notion. In fact, it is a stronger notion:

Lemma 3.8. Multi-leximin is a DSL strategy, but not necessarily vice-versa.

Lemma 3.9. *When there is a finite amount of safety level strategies, multi-leximin is guaranteed to exist.*

An important advantage of the DSL definition is that it naturally extends to settings with continuous outcomes. It is not clear how to extend the leximin definition to such cases. Thus, one possible way of thinking about the DSL notion is that it is a somewhat weaker notion of multi-leximin, that can be used in continuous settings, as well as discrete ones.

4 Main Result: Application to VCG under False-name Attacks

We now move on to present our main result and through it the usefulness of the DSL notion. False-name attacks by an agent i in a combinatorial auction are where instead of sending one combinatorial bid, the agent sends multiple combinatorial bids (a vector \mathbf{b}_i rather than a single bid b_i). The agent then gets all the items allocated to the “agents” (which we call Sybil agents or Sybil bids) $1 \leq j \leq |\mathbf{b}_i|$, and pays the sum of all their payments. Before formally introducing the VCG notations, we note three complexities that are present in our notations: (1) We consider both the notion of DSL strategies (which has the single agent perspective vs. nature states) and social welfare (which accounts for n different agents). (2) We consider welfare for the real n underlying agents of the auction, but since each may use Sybil identities, the VCG allocations are in terms of the Sybil identities. We allow for both by using sub-indexing. (3) Similar to the case of the first-price auction, discretization of the bid space is essential to the result (a counter-example for continuous VCG appears in the full version of our paper. To further simplify the proof, we also assume that the valuation space is discrete, though this assumption can be removed. We allow more granularity to the bid space: valuations are on an ε grid, while bids are on an $\frac{\varepsilon}{2^{|\mathcal{M}|}}$ grid.

Definition 4.1. *Grid(ε) = $\{\varepsilon k\}_{k \in \mathcal{N}} = \{0, \varepsilon, 2\varepsilon, \dots\}$. A combinatorial bid $b \in B$ over an item set M is a function $b : P(M) \rightarrow \text{Grid}(\frac{\varepsilon}{2^{|\mathcal{M}|}})$ from the power set of M to a non-negative bid value. A combinatorial valuation v is similarly $v : P(M) \rightarrow \text{Grid}(\varepsilon)$. With the possibility of Sybil attacks, an agent i with valuation (type) v_i sends a vector of bids (action) $\mathbf{b}_i \in B^*$ (i.e., any amount of combinatorial bids), and faces a nature state $\mathbf{b}_{-i} \in B^*$.*

Let $\eta_i = |\mathbf{b}_i|$, $\eta_{-i} = |\mathbf{b}_{-i}|$ be the number of (Sybil) agents in each vector. An allocation $\alpha^S(\mathbf{b}_i, \mathbf{b}_{-i})$ maps the bid vectors into a partition of S into subsets. We allow indexing $\alpha_1, \dots, \alpha_n$ to mean the union of items allocated to the Sybil identities of each real agent, as well as sub-indexing $\alpha_{i_1}, \dots, \alpha_{i_{\eta_i}}$ to mean the items allocated to a specific Sybil identity of agent i . We denote $SW_\alpha^{Obs} = \sum_{i=1}^n \sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j})$, $SW_\alpha^{Real} = \sum_{i=1}^n v_i(\alpha_i)$ for the observed social welfare of an allocation as specified in the (possibly Sybil) bids, and the real social welfare of the agents, respectively. We denote $truth_i = v_i$ for the truthful bid.

The VCG combinatorial auction is the pair of allocation rule $\alpha^M(\mathbf{b}_i, \mathbf{b}_{-i}) = \arg \max_{\tilde{\alpha}^M(\mathbf{b}_i, \mathbf{b}_{-i})} (SW_{\tilde{\alpha}^M(\mathbf{b}_i, \mathbf{b}_{-i})}^{Obs})$, and the payment rule

$$p_{i_j}^M(\mathbf{b}_i, \mathbf{b}_{-i}) = SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - SW_{\alpha_{M \setminus a_{i_j}^M}^{Obs}}(\mathbf{b}_i, \mathbf{b}_{-i}).$$

Finally, the utility of agent i is $u_i(\mathbf{b}_i, \mathbf{b}_{-i} | v_i) = v_i \left(\bigcup_{1 \leq j \leq \eta_i} a_{i_j}^M(\mathbf{b}_i, \mathbf{b}_{-i}) \right) - \sum_{j=1}^{\eta_i} p_{i_j}^M(\mathbf{b}_i, \mathbf{b}_{-i})$.

Theorem 4.2. *When all bidders play DSL strategies, discrete VCG achieves optimal welfare, even under the possibility of false-name attacks and with general valuations.*

Proof. Our proof follows the following structure: First, we define overbidding Sybil attacks and show that they are not DSL. We then define underbidding attacks and show that they are not DSL. For any of the remaining attacks, which we call exact-bidding (bidding truthfully is also exact-bidding, but not

exclusively so), we show that even though they are not necessarily truthful, they yield maximal welfare. However, this still does not guarantee that one of the remaining strategies is in fact DSL. For this purpose, we show that there exists a DSL strategy: being truthful³.

First, we show that if the Sybil bids are overbidding v (in a sense that will be immediately defined), then, similarly to our proof for the first-price auction (see Appendix B), it is not safety level and thus not DSL. This requires slightly more care since the bids are combinatorial and there are several Sybil bids. We say that $\mathbf{b}_i = (b_{i_1}, \dots, b_{i_{\eta_i}})$ is overbidding if there is a set S and an allocation $\alpha^S(\mathbf{b}_i)$ so that $\sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S(\mathbf{b}_i)) > v_i(S)$.

Claim 4.3. *Overbidding \implies not DSL.*

Proof. Let us choose a maximizing allocation $\alpha^S(\mathbf{b}_i)$ for S .

We denote $\bar{b} = \max_{1 \leq j \leq \eta_i} \max_{S' \subseteq M} b_{i_j}(S') + v_i(S')$ for a number high enough that if some other agent bids it for any subset of M , both the truthful bid v_i or the Sybil attack \mathbf{b}_i will lose that subset. We will use it in our construction of nature states. By the overbidding condition, we can take the average $\tilde{b} = \frac{v_i(S)}{2} + \frac{1}{2} \sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S)$, so that $\sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) > \tilde{b} > v_i(S)$. Consider a nature state where the false-name attacker faces exactly one additive bidder b' that has for any good $g \in M \setminus S$, $b'(g) = \bar{b}$, and for any good $g \in S$, $b'(g) = \frac{\tilde{b}}{|S|}$. The optimal observed welfare allocation is to allocate all goods in $M \setminus S$ to b' , and allocate the set S as in $\alpha_i(S)$. The payment of bidder b_i must be at least $b'(S) = \tilde{b} > v_i(S)$. Therefore, the attacker has negative utility in this case, while truthfulness is individually rational: i.e., it is not a safety level strategy and so also not DSL. \square

We say that $b_{i_1}, b_{i_{\eta_i}}$ are underbidding if there is a set S so that for any allocation $\alpha^S \stackrel{\text{def}}{=} \alpha^S(\mathbf{b}_i)$ so that $\sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) < v_i(S)$.

Claim 4.4. *Underbidding \implies not DSL.*

Proof. Let $\tilde{b} = \frac{1}{2} \sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) + \frac{v_i(S)}{2}$, then

$$\sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S(\mathbf{b}_i)) < \tilde{b} < v_i(S).$$

Let b' be constructed as in the overbidding case. The allocation $\alpha^M(\mathbf{b}_i, b')$ allocates no items to the Sybil bidders of agent i . However, the allocation given agent i bids truthfully $\alpha^M(\text{truth}_i, b')$, allocates the set S to her with payment \tilde{b} , which yields agent i a positive utility $v_i(S) - \tilde{b}$. This yields

$$\min_{\mathbf{b}_{-i} \in D_{v_i}(\mathbf{b}_i, \text{truth}_i)} u_i(\mathbf{b}_i, \mathbf{b}_{-i} | v_i) = 0.$$

On the other hand, we claim that since we know DSL strategies are not overbidding, there are no nature states for which an underbidding Sybil attack gets positive utility while bidding truthfully gets 0 utility. Assume towards contradiction truth_i gets 0 utility. It then either does not win any item, or wins some set

³Another, albeit non-constructive method to show there exists a DSL strategy is by showing the finiteness of undominated exact-bidding Sybil attacks, and then use Corollary 3.5

S and pays $v_i(S)$ for it. Let S be the set that the Sybil bidders win to gain positive utility. As there is no overbidding, this set can be won by $truth_i$ as well (in the respective maximizing allocation)⁴. Then,

$$\begin{aligned}
& SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) && \text{(No overbidding)} \\
& = SW_{\alpha^{M \setminus S}}^{Obs}(truth_i, \mathbf{b}_{-i}) && (i \text{ wins only } S) \\
& = SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_{-i}) && (i\text{'s truthful payment)} \\
& = SW_{\alpha^M}^{Obs}(\mathbf{b}_{-i}) - v_i(S) && \text{(More bids)} \\
& \leq SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - v_i(S)
\end{aligned}$$

So

$$v_i(S) \leq SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) \quad (1)$$

Since our choice of S assumes the Sybil bids win exactly it, we have

$$SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) + SW_{\alpha^S}^{Obs}(\mathbf{b}_i) = SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}),$$

and so, together with Eq. 1,

$$v_i(S) \leq SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) = SW_{\alpha^S}^{Obs}(\mathbf{b}_i).$$

Since there is no overbidding, $SW_{\alpha^S}^{Obs}(\mathbf{b}_i) = v_i(S)$.

We now show that any Sybil bidder j pays $v_i(S) - \sum_{1 \leq t \neq j \leq \eta_i} b_{ij}(\alpha_{ij})$. Since S is allocated to the Sybil bidders and $M \setminus S$ to others,

$$\begin{aligned}
& SW_{\alpha^{M \setminus a_{ij}^M}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) = \\
& SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) + SW_{\alpha^{S \setminus a_{ij}^M}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i})
\end{aligned} \quad (2)$$

Then,

$$\begin{aligned}
p_{ij}^M & = SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - SW_{\alpha^{M \setminus a_{ij}^M}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) \\
& = SW_{\alpha^M}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) - SW_{\alpha^{M \setminus S}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) \\
& \quad - SW_{\alpha^{S \setminus a_{ij}^M}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) \\
& = v_i(S) - SW_{\alpha^{S \setminus a_{ij}^M}}^{Obs}(\mathbf{b}_i, \mathbf{b}_{-i}) \\
& = v_i(S) - \sum_{j=1}^{\eta_i} b_{ij}(\alpha_{ij})
\end{aligned}$$

The total payment of agent i is then

$$\begin{aligned}
\sum_{j=1}^{\eta_i} p_{ij}^M & = \sum_{j=1}^{\eta_i} v_i(S) - \sum_{1 \leq t \neq j \leq \eta_i} b_{ij}(\alpha_{ij}) = \\
\eta_i \cdot v_i(S) & - (\eta_i - 1) \sum_{j=1}^{\eta_i} b_{ij}(\alpha_{ij}) = v_i(S).
\end{aligned}$$

⁴In full generality, $truth_i$ may win a set s that has partial intersection with S . The analysis of this case is essentially the same, and stems from the fact that the alternative value for the items forces zero utility on the truthful agent.

This concludes that whenever the utility of $truth_i$ is 0, then the utility for the Sybil attack is 0 as well.

In any other case, the utility of $truth_i$ must be strictly positive, and since the bids are discrete the minimum over all these cases satisfies

$$\min_{\mathbf{b}_{-i} \in D_{v_i}(\mathbf{b}_i, truth_i)} u_i(truth_i, \mathbf{b}_{-i} | v_i) \geq \frac{1}{2|M|!}.$$

Therefore, underbidding is not DSL. \square

We consider *exact-bidding* such Sybil bids that have for any set of items S , $\max_{\alpha^S(\mathbf{b}_i)} \sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S(S)) = v_i(S)$. These are exactly all the Sybil attacks that are neither overbidding nor underbidding. $truth_i$ is also exact-bidding.

Claim 4.5. *Exact-bidding \implies optimal welfare.*

Proof. Consider an allocation $\alpha_F \stackrel{def}{=} \alpha^M(\mathbf{b}_i, \mathbf{b}_{-i})$ attained when all players choose an exact-bidding attack, vs $\alpha_T \stackrel{def}{=} \alpha^M(truth_i, \mathbf{truth}_{-i})$. We have

$$\begin{aligned} SW_{\alpha_T}^{Real} &\leq && \text{(Truthful)} \\ SW_{\alpha_T}^{Obs} &\leq && \text{(No underbidding)} \\ SW_{\alpha_F}^{Obs} &\leq && \text{(No overbidding)} \\ SW_{\alpha_F}^{Real} &&& \end{aligned}$$

In words, since there is no underbidding in the Sybil attack, if we take the set allocated to each agent i under the allocation that maximizes welfare under truthfulness, there are Sybil bidders i_{j_1}, \dots, i_{j_k} with the same aggregate valuation for it. So, $SW_{\alpha_F}^{Obs}$ is lower bounded by the optimal truthful welfare. Since there is also no overbidding, whatever allocation is chosen as α_F^{Obs} is at least as good to each agent i as is declared. \square

Claim 4.6. *$truth_i$ is DSL.*

Proof. Consider some exact-bidding Sybil attack \mathbf{b}_i .

Case 1: There is a set S so that $\forall 1 \leq j \leq \eta_i, b_{i_j}(S) < v_i(S)$. Then, by the exact-bidding condition there must be some allocation $\alpha^S(\mathbf{b}_i)$ (with at least two non-empty allocations $\alpha_{i_j}^S$) so that $\max_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) < \sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) = v_i(S)$. Consider the nature state where there is one bid b' so that $b'(\alpha_{i_j}^S) = v_i(\alpha_{i_j}^S)$ for any $1 \leq j \leq \eta_i$, and the rest of the sets are defined upward-monotonely: They inherit the largest value of a subset. With this nature state, the Sybil attack has utility 0. On the other hand, $truth_i$ has positive utility of $v_i(S) - \max_{j=1}^{\eta_i} v_i(\alpha_{i_j}^S) > 0$. Since $truth_i$ is individually rational, it is thus DSL w.r.t. such Sybil attacks.

Case 2: For every set S , there is such j' with $b_{i_{j'}}(S) = v_i(S)$. It must hold by the exact-bidding condition that for any allocation $\alpha^S(\mathbf{b}_i)$, $\sum_{j=1}^{\eta_i} b_{i_j}(\alpha_{i_j}^S) \leq v_i(S) = b_{i_{j'}}(S)$. We may assume that VCG prefers to assign larger bundles when tie-breaking between possible assignments. Then, it must be that any allocation to the Sybil bidders is given to one Sybil bidder as a whole bundle. It is then weakly better to send only $b_{i_{j'}}$ as a single bid instead of \mathbf{b}_i . Furthermore, it is then weakly better to send $truth_i$, since truthfulness is dominant for single bid VCG. Since this is true given any nature state, the Sybil attack is weakly dominated by $truth_i$, which implies $truth_i$ is DSL with respect to it.

This covers all the exact-bidding Sybil attacks. DSL strategies with respect to overbidding and underbidding attacks are implied by the relevant discussion. Overall this covers all Sybil attacks. \square



5 Discussion and Future Directions

In the example of the discrete first-price auction in Section 2, as well as in our main result in Section 4, the DSL solution concept leads to optimal results: truthfulness (or near truthfulness), and optimal revenue or welfare. In Appendix B we study the first-price auction, and show similar results for its discrete variant. However, for the classic setting of voting, we show in the full version of this paper that this is not the case, and that solutions may have various surprising forms.

A robust notion missing from our discussion in Section 3 is min-max regret. We show in Appendix A it does not imply or is implied by our notion of DSL, and give further characteristics of it. It is also compared with our notion as part of our discussion of the discrete first-price auction in Appendix B.

In our definition of DSL, we consider only pure nature states. We justify this choice in Appendix C of the full version, by showing that if we consider mixed nature states as well, then the DSL and safety level notions become one. In Appendix D of the full version, we show a possible refinement of our notion of DSL, and demonstrate why it may be useful.

A few immediate open questions follow our work:

- We find that DSL is a stronger notion than safety level. In settings previously studied that proved performance guarantees for safety level strategies, do DSL strategies exist? Can they yield better performance guarantees?
- In the case of single-item auctions, our analysis of the discrete first price auction implies that with DSL bidders, it is possible to achieve optimal welfare and revenue. Does this extend to combinatorial auctions? If so, does it hold even when the discretization must be polynomially bound?
- In the presence of partial knowledge or the option to elicitate it (similar to the ideas in [14]), what would the DSL action be? This is relevant, for example, when agents arrive sequentially, and so the set of feasible nature states diminishes for later agents.

Acknowledgements

Yotam Gafni and Moshe Tennenholtz were supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant No. 740435).

References

- [1] Colleen Alkalay-Houlihan & Adrian Vetta (2014): *False-Name Bidding and Economic Efficiency in Combinatorial Auctions*. *Proceedings of the AAAI Conference on Artificial Intelligence* 28(1), doi:10.1609/aaai.v28i1.8828. Available at <https://ojs.aaai.org/index.php/AAAI/article/view/8828>.
- [2] Robert J Aumann (1985): *On the non-transferable utility value: A comment on the Roth-Shafer examples*. *Econometrica: Journal of the Econometric Society*, pp. 667–677, doi:10.2307/1911661.
- [3] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang & Ariel D. Procaccia (2016): *Handbook of Computational Social Choice*, 1st edition. Cambridge University Press, USA, doi:10.1017/CBO9781107446984.
- [4] George Christodoulou, Annamária Kovács & Michael Schapira (2016): *Bayesian Combinatorial Auctions*. *Journal of the ACM* 63(2), doi:10.1145/2835172. Available at <https://doi.org/10.1145/2835172>.

- [5] Michael Suk-Young Chwe (1989): *The discrete bid first auction*. *Economics Letters* 31(4), pp. 303–306, doi:[https://doi.org/10.1016/0165-1765\(89\)90019-0](https://doi.org/10.1016/0165-1765(89)90019-0). Available at <https://www.sciencedirect.com/science/article/pii/0165176589900190>.
- [6] Sofie De Clercq, Steven Schockaert, Ann Nowé & Martine De Cock (2018): *Modelling incomplete information in Boolean games using possibilistic logic*. *International Journal of Approximate Reasoning* 93, pp. 1–23, doi:<https://doi.org/10.1016/j.ijar.2017.10.017>. Available at <https://www.sciencedirect.com/science/article/pii/S0888613X17302554>.
- [7] Didier Dubois & Philippe Fortemps (1999): *Computing improved optimal solutions to max–min flexible constraint satisfaction problems*. *European Journal of Operational Research* 118(1), pp. 95–126, doi:[https://doi.org/10.1016/S0377-2217\(98\)00307-5](https://doi.org/10.1016/S0377-2217(98)00307-5). Available at <https://www.sciencedirect.com/science/article/pii/S0377221798003075>.
- [8] Michal Feldman, Hu Fu, Nick Gravin & Brendan Lucier (2013): *Simultaneous Auctions Are (Almost) Efficient*. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC '13*, Association for Computing Machinery, New York, NY, USA, p. 201–210, doi:10.1145/2488608.2488634. Available at <https://doi.org/10.1145/2488608.2488634>.
- [9] John A. Ferejohn & Morris P. Fiorina (1974): *The Paradox of Not Voting: A Decision Theoretic Analysis*. *American Political Science Review* 68(2), p. 525–536, doi:10.2307/1959502.
- [10] Yotam Gafni, Ron Lavi & Moshe Tennenholtz (2020): *VCG under Sybil (False-Name) Attacks - A Bayesian Analysis*. *Proceedings of the AAI Conference on Artificial Intelligence* 34(02), pp. 1966–1973, doi:10.1609/aaai.v34i02.5567. Available at <https://ojs.aaai.org/index.php/AAAI/article/view/5567>.
- [11] Paul Harrenstein, Wiebe van der Hoek, John-Jules Meyer & Cees Witteveen (2001): *Boolean games*. In: *Proceedings of the 8th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 287–298.
- [12] Atsushi Iwasaki, Vincent Conitzer, Yoshifusa Omori, Yuko Sakurai, Taiki Todo, Mingyu Guo & Makoto Yokoo (2010): *Worst-Case Efficiency Ratio in False-Name-Proof Combinatorial Auction Mechanisms*. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, p. 633–640.
- [13] Benny Lehmann, Daniel Lehmann & Noam Nisan (2006): *Combinatorial auctions with decreasing marginal utilities*. *Games and Economic Behavior* 55(2), pp. 270–296, doi:10.1016/j.geb.2005.02.006.
- [14] Tyler Lu & Craig Boutilier (2011): *Robust Approximation and Incremental Elicitation in Voting Protocols*. In: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume One*, IJCAI'11, AAAI Press, p. 287–293, doi:10.5591/978-1-57735-516-8/IJCAI11-058.
- [15] S. Merrill (1982): *Strategic Voting in Multicandidate Elections under Uncertainty and under Risk*. In Manfred J. Holler, editor: *Power, Voting, and Voting Power*, Physica-Verlag HD, Heidelberg, pp. 179–187.
- [16] Hervé Moulin (2003): *Fair Division and Collective Welfare*. The MIT Press, doi:10.7551/mitpress/2954.001.0001. Available at <https://doi.org/10.7551/mitpress/2954.001.0001>.
- [17] Noam Nisan, Tim Roughgarden, Éva Tardos & Vijay V. Vazirani (2007): *Algorithmic Game Theory*. Cambridge University Press, doi:10.1017/CBO9780511800481.
- [18] L. J. Savage (1951): *The Theory of Statistical Decision*. *Journal of the American Statistical Association* 46(253), pp. 55–67, doi:10.1080/01621459.1951.10500768. arXiv:<https://www.tandfonline.com/doi/pdf/10.1080/01621459.1951.10500768>.
- [19] A. Sen (1970): *Collective Choice and Social Welfare*. Mathematical economics texts, Holden-Day. Available at <https://books.google.co.il/books?id=kzq6AAAAIAAJ>.
- [20] Moshe Tennenholtz (2001): *Rational Competitive Analysis*. In: *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'01, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, p. 1067–1072.

- [21] Moshe Tennenholtz (2002): *Competitive safety analysis: Robust decision-making in multi-agent systems*. *Journal of Artificial Intelligence Research* 17, pp. 363–378, doi:10.1613/jair.1065.
- [22] Makoto Yokoo, Yuko Sakurai & Shigeo Matsubara (2004): *The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions*. *Games and Economic Behavior* 46(1), pp. 174–188, doi:10.1016/S0899-8256(03)00045-9.

A Min-max Regret

Another robust solution notion is the min-max regret [18]. The notion has many uses in voting: [9] showed it can be used to explain why voters choose to participate in elections, and [15] used it to “resolve” the Gibbard-Satterthwaite impossibility theorem (see, e.g., [3]), by showing that plurality voting (for example) is truthful under this notion. [14] showed how when only partial preferences are known, voting rules can use this notion to decide a winner, and design good elicitation schemes.

Definition A.1. *Regret for an action a_i and nature state a_{-i} given a type θ_i is*

$$\text{Reg}(a_i, a_{-i} | \theta_i) = \max_{a'_i} u(a'_i, a_{-i} | \theta_i) - u(a_i, a_{-i} | \theta_i).$$

Max regret for an action a_i given a type θ_i is

$$\text{Reg}(a_i, a_{-i} | \theta_i).$$

A min-max regret action belongs to

$$\arg \min_{a_i} \max_{a_{-i}, a'_i} u(a'_i, a_{-i} | \theta_i) - u(a_i, a_{-i} | \theta_i).$$

In words, the regret of an action a_i under nature state a_{-i} is the maximal lost utility $u(a'_i, a_{-i} | \theta_i) - u(a_i, a_{-i} | \theta_i)$ of choosing a_i instead of a'_i , over all possible actions a'_i (this regret is non-negative, as there is always the option of choosing a_i itself). Max regret is the maximal such regret over all nature states, and the min-max regret action is the action a_i that has minimal max regret.

Proposition A.2. *Dominant strategy \implies min-max regret*

Proof. Consider a dominant strategy s , fix a type θ_i , and let $a = s(\theta_i)$. For any a', a_{-i} , we have that $u(a', a_{-i}) - u(a, a_{-i}) \leq 0$, i.e., the max regret for a is 0, the minimum possible, and so a is a min-max regret action. Since this holds for all types, s is a min-max regret strategy. \square

Example A.3. *Min-max regret $\not\Rightarrow$ safety level*

Consider two players, with actions a, b , and A, B respectively. Consider $u_1(a, A) = u_1(a, B) = 0, u_1(b, A) = -1, u_1(b, B) = 100$. The max regret of a for player 1 is 100, and the max regret of b is 1, and so b is the min-max regret strategy, while a is the unique safety level strategy.

B DSL Strategies: Application to the First-price Auction

B.1 The First-price Auction

As an illustrative example, we demonstrate the usage of our solution concept using the first price and discrete first price single item auctions. Interestingly, we show that in the first-price auction, there are no DSL strategies. However, moving to a discrete setting, we show that in the discrete first-price auction [5], there is a unique DSL strategy, which achieves maximal welfare and near maximal revenue.

Definition B.1. An agent i has a value (type) v_i for an item. The agent's bid b_i (action) and nature states b_{-i} are from the same bid space. The auctioneer allocates the item to the highest bidder (either the agent or nature, tie-breaking towards nature) and if the agent wins it receives $v_i - b_i$, and otherwise 0.

First-price auction (FPA): bid space is $0 \leq b_i \leq v_i$.

Discrete first-price auction (DFPA): bid space is $b_i \in \{\varepsilon \cdot k \mid \varepsilon \cdot k \leq v_i\}_{k \in \mathcal{N}}$.

Note that we reformulate the auctions to suits our agent perspective formulation. Moreover, we omit strategies that are not individually rational (in the (discrete) first-price auction, overbidding has negative utility in some nature states), which is justified by our later discussion in Proposition 3.4. We also ignore multitude in nature states that does not change the auction outcome. I.e., we only consider the highest bids by others as the nature state, and not the entire bid vector. For the DFPA, we denote $\varepsilon_{net}(v_i) = \varepsilon \cdot \max_{\varepsilon n \leq v_i} n$, i.e., the closest possible bid below the agent's value of the item.

In the first-price auction, the notion of DSL strategies is not of much help:

Lemma B.2. *In the first-price auction, there are no DSL bid strategies.*

Proof. First, consider some bid $0 \leq b_i < v_i$. Compare it with another bid b'_i that satisfies $b_i < b'_i < v_i$.⁵ Consider a nature state b_{-i} so that $b_i < b_{-i} < b'_i$. Then, $0 = u_i(b_i, b_{-i} | v_i) \neq u_i(b'_i, b_{-i} | v_i) = v_i - b'_i$. Thus,

$$\min_{b_{-i} \in D_{v_i}(b_i, b'_i)} u_i(b_i, b_{-i} | v_i) = 0.$$

On the other hand, for the bid b'_i and for some nature state b_{-i} , $u_i(b'_i, b_{-i}) = 0$ if and only if $b_{-i} \geq b'_i$. In all such cases, it also holds that $u_i(b_i, b_{-i} | v_i) = 0$. In all other cases, i.e., when $b_{-i} < b'_i$, the utility of the bidder satisfies $u_i(b'_i, b_{-i} | v_i) = v_i - b'_i$. We conclude that

$$\begin{aligned} \min_{b_{-i} \in D_{v_i}(b_i, b'_i)} u_i(b'_i, b_{-i} | v_i) &= v_i - b'_i \\ &> \min_{b_{-i} \in D_{v_i}(b_i, b'_i)} u_i(b_i, b_{-i} | v_i) &= 0, \end{aligned}$$

and the bid strategy b_i is not DSL.

If $b_i = v_i$, then for any nature state b_{-i} , $u_i(b_i, b_{-i} | v_i) = 0$. For some $0 \leq b'_i < b_i$, for any nature state b_{-i} where its utility is non-zero, we have $u_i(b'_i, b_{-i} | v_i) = v_i - b'_i > 0$, and so similarly to before $b_i = v_i$ is not DSL. \square

However, things get more interesting with the DFPA:

Lemma B.3. *In the discrete first-price auction:*

For types that have $\varepsilon_{net}(v_i) \neq v_i$, and types with $\varepsilon_{net}(v_i) = v_i = 0$, bidding $\varepsilon_{net}(v_i)$ is the unique DSL bid.

For types with $\varepsilon_{net}(v_i) = v_i \neq 0$, the unique DSL bid is $\varepsilon_{net}(v_i) - \varepsilon$.

We first give a proof for the pure DSL case.

Proof. The argument why any other bid strategy is not DSL follows a discretized version of the proof for Lemma B.2.

Case 1: $\varepsilon_{net}(v_i) \neq v_i$

⁵Note that the proof is written for the pure DSL case. However, it immediately generalizes to the mixed case, by adapting “ $b_i < v_i$ ” to “has a positive probability to satisfy $b_i < v_i$ ”, etc.

Consider some bid with $0 \leq b'_i < \varepsilon_{net}(v_i)$. By the same argument as in the first part of the proof of Lemma B.2, bidding $\varepsilon_{net}(v_i)$ is DSL w.r.t. b'_i . Since there are no bids with $\varepsilon_{net}(v_i) < b'_i \leq v_i$ by the definition of ε_{net} , we conclude that $\varepsilon_{net}(v_i)$ is DSL w.r.t. all other bids, i.e., DSL.

Case 2: $\varepsilon_{net}(v_i) = v_i = 0$

The unique safety level bid is to bid 0, and so by Proposition 3.4 it is also the unique DSL strategy.

Case 3: $\varepsilon_{net}(v_i) = v_i \neq 0$

Similar to the first case, with the difference that bidding v_i always leads to utility 0, and so the DSL bid bracket is $v_i - \varepsilon$. \square

The following lemma completes the mixed DSL case:

Lemma B.4. *In the discrete first-price auction with mixed strategies, following the unique DSL pure strategy is the unique DSL strategy.*

Proof. We show the proof for case 1 where $\varepsilon_{net}(v_i) \neq v_i$. The other cases are done similarly.

Let s_i be the stated strategy, v_i the valuation (type) and the bid $b = s_i(\theta_i)$. Let b' be some other bid: since $b \neq b'$, the bracket $\varepsilon_{net}(v_i)$ has probability $p < 1$ of being the actualized bid. Consider the case b_{-i} where another bidder bids $\varepsilon_{net}(v) - \varepsilon$, and ties are broken in favor of the other bidder. Then, $u_i(b', b_{-i}|v_i) = \mathbb{E}_{\tilde{b}' \sim b'}[u_i(\tilde{b}', b_{-i}|v_i)] = p \cdot (v_i - \varepsilon_{net}(v_i)) + (1 - p)\mathbb{1}[\tilde{b}' > \varepsilon_{net}(v_i)] \cdot (v_i - \tilde{b}') = p \cdot (v_i - \varepsilon_{net}(v_i)) + (1 - p)\mathbb{1}[\tilde{b}' > v_i](v_i - \tilde{b}') < p \cdot (v_i - \varepsilon_{net}(v_i)) < v_i - \varepsilon_{net}(v_i)$. In any nature state and actualized outcome over the mixed bid b' , if b does not win the item, then b' does not win the item, or, alternatively, it wins it and receives negative utility. So, $\min_{b_{-i} \in D_{v_i}(b, b')} u_i(b, b_{-i}|v_i) \geq v_i - \varepsilon_{net}(v_i) > \min_{b_{-i} \in D_{v_i}(b, b')} u_i(b', b_{-i}|v_i)$, and so by the DSL condition b' is not DSL (and b is DSL w.r.t. b'). \square

The simple intuition as to why the discrete first-price auction “works” (to guarantee a DSL strategy) and the first-price auction does not, is that in the first-price auction there is always a “safer” bid that would guarantee winning the item in more nature states. In the discrete first-price auction, due to bracketing, the highest bracket that can have positive utility is that DSL bid. Note that this is “almost” truthful: When $\varepsilon_{net}(v_i) \neq v_i$, it is the closest bracket to v_i , and it is less than ε away from it. When $\varepsilon_{net}(v_i) = v_i$ (which should be seen as a rare case, where the value precisely matches the epsilon net), it is not the truthful bracket, but it is ε close to it. It is also very close to optimal revenue for the auctioneer: If n individually rational agents participate, the most the auctioneer can get is $\max_{1 \leq i \leq n} v_i$. If they play DSL strategies, she will get at least $\max_{1 \leq i \leq n} v_i - \varepsilon$.

We note that for the discrete first-price auction, DSL identifies with multi-leximin.

Corollary B.5. *The unique DSL strategy of the discrete first-price auction is also the unique multi-leximin strategy.*

Proof. For an agent i with value v_i there is a finite amount of safety level strategies, namely all the strategies with $b_i \leq v_i$, the amount of which is at most $\lceil \frac{v_i}{\varepsilon} \rceil + 1$. By Lemma 3.9, there must exist a multi-leximin strategy. By Lemma 3.8 it is also DSL. Since there is a unique DSL strategy by Lemma B.3, it must also be the unique multi-leximin. \square

On the other hand, we now see that min-max regret yields a different solution to the discrete first-price auction than DSL, i.e., the two notions do not imply each other. [20] previously applied min-max regret in auction settings, and in particular discussed the DFPA in their Claim 3.1, which we restate adapted to our notations:

Claim B.6. *In the discrete first-price auction, the min-max regret strategy is to bid $\varepsilon_{net}(\frac{v_i}{2})$.*

Proof. For any bid b_i , the maximum regret is either b_i itself (in the case when no other bidders show up and it was possible to bid and pay 0), or $v_i - (b_i + \varepsilon)$ (in the case when another bidder bids b_i and the item goes to her.⁶ We are thus looking for $\arg \min_{b_i} \max\{b_i, v_i - b_i - \varepsilon\}$, among the ε_{net} feasible bids.

For $b'_i < \varepsilon_{net}(\frac{v_i}{2})$, the regret is thus at least

$$\begin{aligned} \text{Reg}(b'_i) &\geq v_i - b_i - \varepsilon \\ &\geq v_i - (\varepsilon_{net}(\frac{v_i}{2}) - \varepsilon) - \varepsilon = v_i - \varepsilon_{net}(\frac{v_i}{2}) \\ &\geq \max\{v_i - \varepsilon_{net}(\frac{v_i}{2}) - \varepsilon, \frac{v_i}{2}\} \\ &\geq \max\{v_i - \varepsilon_{net}(\frac{v_i}{2}) - \varepsilon, \varepsilon_{net}(\frac{v_i}{2})\} = \text{Reg}(b_i). \end{aligned}$$

For $b'_i > \varepsilon_{net}(\frac{v_i}{2})$, the regret is at least

$$\begin{aligned} \text{Reg}(b'_i) &\geq b'_i \geq \varepsilon_{net}(\frac{v_i}{2}) + \varepsilon \\ &\geq \max\{\varepsilon_{net}(\frac{v_i}{2}), \frac{v_i}{2}\} \\ &\geq \max\{\varepsilon_{net}(\frac{v_i}{2}), v_i - \varepsilon_{net}(\frac{v_i}{2}) - \varepsilon\} = \text{Reg}(b_i). \end{aligned}$$

We conclude that $\varepsilon_{net}(\frac{v_i}{2})$ is the min-max regret bid strategy. □

Finally, we use the discrete first-price auction to demonstrate the difference between leximin and DSL strategies.

Example B.7. *We demonstrate that leximin is different from DSL using the discrete first-price auction. Bidding 0 is the leximin action, as its set of outcomes is simply the set of two items $U_0 = \{0, v_i\}$: This is the leximin since any other bid $b_i > 0$ has $U_{b_i} = \{0, v_i - b_i\}$.*

⁶This is true under worst-case arbitrary tie-breaking. If tie-breaking is uniformly random between bidders of the same bracket, this is still true as the limiting regret when there are $n \rightarrow \infty$ bidders in the same bracket