

Stability Scores: Measuring Coalitional Stability

Michal Feldman
Harvard University and Hebrew
University of Jerusalem*
mfeldman@seas.harvard.edu

Reshef Meir
Microsoft Research and Hebrew
University of Jerusalem
reshef.meir@mail.huji.ac.il

Moshe Tennenholtz
Microsoft Research and
Technion-Israel Institute of
Technology
moshet@microsoft.com

ABSTRACT

We introduce a measure for the level of stability against coalitional deviations, called *stability scores*, which generalizes widely used notions of stability in non-cooperative games. We use the proposed measure to compare various Nash equilibria in congestion games, and to quantify the effect of game parameters on coalitional stability. For our main results, we apply stability scores to analyze and compare the Generalized Second Price (GSP) and Vickrey-Clarke-Groves (VCG) ad auctions. We show that while a central result of the ad auctions literature is that the GSP and VCG auctions implement the same outcome in one of the equilibria of GSP, the GSP outcome is far more stable. Finally, a modified version of VCG is introduced, which is group strategy-proof, and thereby achieves the highest possible stability score.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*;

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Theory, Economics

Keywords

Game theory, Auctions, Mechanism design, Solution concepts

1. INTRODUCTION

One of the most basic questions of game theory is: given a game in strategic form, what is its solution? By *solution* we typically mean a strategy profile that can be proposed to all agents, and no rational agent would want to deviate from it. Thus a solution should be *stable*. Many solution concepts for games have been studied; these studies differ by the level and interpretation of stability, as well as by the underlying assumptions that are required to achieve it. The best known solution concept for games is the Nash equilibrium (NE), a strategy profile from which no agent has an incentive to deviate *unilaterally*.

*At the time of research the author was affiliated with Microsoft Research in Herzlia, Israel.

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

A basic problem with the NE solution concept is that a NE does not take into account joint deviations by coalitions of players. We usually assume that an individual will deviate from a profile if she has an available strategy that strictly increases her payoff. In some settings it would be natural to assume also that a group of individuals will deviate if they have an available joint strategy that strictly increases the payoff of each group member. The *Strong Equilibrium* (SE) concept by Aumann [2] deals with this problem. A profile is a SE if no coalition of agents can jointly deviate in a way that strictly increases the payoff of each coalition member. Intermediate levels of coalitional stability have been suggested, such as stability against deviations of small coalitions (see e.g. [1]), and in particular pairs. An even more appealing solution concept than SE is the *Super-Strong Equilibrium* (SSE) that considers deviations in which no member loses while at least one member makes a positive gain (see, for example, [11]).

A major problem with these proposed solutions is that they seldom exist. Indeed, SSE rarely exist even in cases where strong equilibria do exist (e.g., in simple congestion games [13, 1]), and even if only deviations by pairs are considered.

In this paper we relax the strong requirement that no coalition will have an incentive to deviate, and suggest a quantitative measure to coalitional stability. Assuming we have a Nash equilibrium profile of a game where some pairs of agents can still deviate, we may still wish to measure its stability by referring to the *number* of pairs that have beneficial deviations from that profile. More generally, given a game and a strategy profile, we can associate with it a tuple in which the r -th entry in the tuple is the number of coalitions of size r that can gain by a deviation. This tuple determines the *stability score* of the strategy profile.

Given two strategy profiles, we need a way to decide which one is more stable. A common practice in game theory is to prefer strategy profiles that are *in equilibrium*, i.e. in which there are no unilateral deviations. Since small coalitions are more likely to form and maintain cooperation, a natural extension is to compare stability scores of games with associated strategy profiles using a *lexicographic* ordering of the corresponding vectors.¹ For example, given two n -person games, G_1 and G_2 , with respective Nash equilibria s_1 and s_2 , the stability score of the former will be higher if the number of beneficial deviations by pairs from s_1 in G_1 is smaller than the number of beneficial deviations by pairs from s_2 in G_2 .

While the existence of, say, 19 coalitions that can deviate rather than 15 does not have much significance, we usually care about the behavior in some parameterized family of games where parameters

¹There are many ways to compare stability score vectors. Choosing the “right” one highly depends on the context and underlying assumptions. However in this paper we avoid such complications by only comparing deviations of coalitions of the same size.

may include number of players, size of the strategy space, etc. If the score of \mathbf{a} is *asymptotically lower* than the score of \mathbf{a}' (w.r.t. one of the parameters), then this may indicate that \mathbf{a}' is substantially more prone to coalitional deviations.

Moreover, when studying such a parametrized family, stability scores may assist us in understanding how the parameters of the game affect coalitional stability. This holds even if there is a unique or a prominent equilibrium.

Stability scores are particularly useful in the context of *mechanism design*, as they allow us to quantify the coalitional stability of various mechanisms and to compare mechanisms that operate in a specific domain. To illustrate this point, we consider two central mechanisms in what is perhaps the most widely studied economic setup in recent years: ad auctions. We analyze in detail the Generalized Second Price (GSP) auction and the Vickrey-Clarke-Groves (VCG) auction, and compare their stability scores.

1.1 Related work

Related solution concepts in games

In the context of non-cooperative games approximate stability is typically measured by the strength of the incentive required to convince an agent to deviate, captured for example by the concept of ϵ -Nash equilibrium. As discussed above, stability against collusion is captured by concepts such as SE and SSE, but these often do not allow a fine distinction between various outcomes.

In addition, coalitions are the key component in *cooperative* game theory, and many variations of coalitional stability have been studied. While we are unaware of solutions concepts that quantify stability by measuring coalitional deviations, models of restricted cooperation capture social constraints that may prevent the formation of some coalitions [18]. Thus a (cooperative) game may not be stable against every coalitional deviation (i.e. have an empty core), but still satisfy all the coalitions that can form in practice. Recently, some papers studied how such social context affects the stability of the game [7, 17]. Moreover, even if some coalition *can* gain by deviation, it may or may not do so: Members of the coalition might intentionally avoid cooperation based on far-sighted prediction (an assumption underlying *coalition-proofness* for example [4]), or just fail to recognize the benefit in deviating. This is especially true if the coalition is large. Stability scores do not assume a particular social context or incentive structure, but simply try and minimize the number of coalitions with profitable deviations.

Collusion and equilibria in ad auctions

Major results of previous work on ad auctions, characterized a special family of equilibria of GSP the auction (used in practice), termed *Symmetric Nash Equilibria*, or SNE (see Section 4.1 for details) [20]. SNEs have many attractive properties which make them a natural choice as outcomes of the GSP auction. Moreover, it has been shown that the SNE leading to the lowest revenue for the seller (termed *Lower Equilibrium* (LE)), coincides with the natural equilibrium of VCG where all bidders report their true values.

The above results led to a surge of papers comparing VCG and the various equilibrium outcomes of GSP, under both public information and private information settings [14, 19, 9, 15]. However, these comparisons focused mainly on revenue, rather than on coalitional stability. The VCG mechanism was shown to be vulnerable to collusion in various domains (see, e.g., [6, 3] for relatively recent work), compared to a simple first-price (pay-your-bid) auction. The formal literature on collusion in second-price auctions goes back to Graham and Marshall [12], while the literature on the more involved matter of collusion in first-price auctions goes back

to McAfee and McMillan [16].

1.2 Our contribution

Stability scores are formally defined in Section 2, where we show how they generalize well known solution concepts. In Section 3 we study strict stability scores in a simple family of congestion games. The main purpose of this study is to demonstrate how stability scores can be used in order to compare different Nash equilibria, and to measure how stability is affected by game's parameters. Moreover, while the studied family itself is quite simple, it is often used to model real world situations such as load balancing. Our analysis can give some intuition as to the main factors affecting coalitional stability in such games.

The main results are in Section 4, where we present the VCG and GSP mechanisms for ad auctions (adopting the original model advocated for that setting in the seminal work by Varian [20] and by Edelman et al. [8]), and show bounds on stability scores in these auctions. In particular, we study how the stability of GSP varies as a function of the distributions of agents' valuations and slots' click-through rates, thereby showing that under certain reasonable conditions GSP is far more stable than VCG.

In Section 5 we introduce a modification to the VCG auction that can be used to overcome the observed instability of VCG. In particular, we show that a revised VCG, in which a random reserve price is introduced, induces truth-telling as a super-strong equilibrium.

Omitted proofs are available in the full version of this paper [10].

2. PRELIMINARIES

Games and equilibria

Let $G = \langle N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a normal form game, where $N = \{1, \dots, n\}$ is the set of players, A_i is the set of actions available to player i , and $u_i : \mathbf{A} \rightarrow \mathbb{R}$ is player i 's utility, where $\mathbf{A} = A_1 \times \dots \times A_n$ is the set of joint actions (profiles), and for every $\mathbf{a} \in \mathbf{A}$, $u_i(\mathbf{a})$ denotes the utility of player i under action profile \mathbf{a} . The vector of actions of all players except player i in the profile \mathbf{a} is denoted by a_{-i} . An action profile $\mathbf{a} \in \mathbf{A}$ is a *Nash Equilibrium* (NE) if $u_i(\mathbf{a}) \geq u_i(b_i, a_{-i})$ for every agent $i \in N$ and every alternative action $b_i \in A_i$.

When considering coalitions, given an action profile \mathbf{a} , we denote by a_S the profile of agents in S , and by A_S the set of all such joint actions. The profile of all agents in $N \setminus S$ is denoted by a_{-S} .

Given a profile of actions $\mathbf{a} \in \mathbf{A}$, $b_S \in A_S$ is a *strict deviation* from \mathbf{a} if $u_i(b_S, a_{-S}) > u_i(a_S, a_{-S})$ for every $i \in S$. The profile \mathbf{a} is termed a *Strong Equilibrium* (SE) if there are no $S \subseteq N$ and $b_S \in A_S$, such that b_S is a strict deviation from \mathbf{a} .

One can also consider the following weaker notion of deviation. Given a profile of actions $\mathbf{a} \in \mathbf{A}$, $b_S \in A_S$ is a *deviation* from \mathbf{a} if $u_i(b_S, a_{-S}) \geq u_i(a_S, a_{-S})$ for every $i \in S$ and there exists $j \in S$ such that $u_j(b_S, a_{-S}) > u_j(a_S, a_{-S})$. The profile \mathbf{a} is termed a *Super-Strong Equilibrium* (SSE) if there is no $S \subseteq N$, $b_S \in A_S$ that is a deviation from \mathbf{a} . Since every strict deviation is clearly a deviation, every SSE is also a SE.

SSE captures the natural requirement that we should resist even situations in which a deviation only benefits some of the deviators without hurting others. A strategy profile is *r-SE* (respectively, *r-SSE*) if there are no coalitions of size at most r that have strict deviations (resp., deviations).

Stability scores

The stability score of the profile \mathbf{a} in game G is defined as a vector with n entries. For every $1 \leq r \leq n$, let $\mathcal{D}_r(G, \mathbf{a}) \in \mathbb{N}$ (respectively, $\mathcal{SD}_r(G, \mathbf{a}) \in \mathbb{N}$) be the *number* of coalitions of size r that

have deviations (resp., strict deviations) from \mathbf{a} in G . While there are many ways to impose an order on equilibria based on these vectors, we believe that the following lexicographic order is particularly natural.

Given two n -player games G and G' and two profiles \mathbf{a} and \mathbf{a}' in the respective games, we say that the pair (G, \mathbf{a}) is *more resistant to deviations* (or *more stable*) than (G', \mathbf{a}') , if there exists some $r \leq n$ such that $\mathcal{D}_r(G, \mathbf{a}) < \mathcal{D}_r(G', \mathbf{a}')$ and the terms are equal for every $r' < r$. We can similarly compare strict stability scores to one another.

Our definition of stability score generalizes some widely used notions of stability. For example, \mathbf{a} is a Nash equilibrium (NE) of G iff $\mathcal{D}_1(G, \mathbf{a}) = \mathcal{SD}_1(G, \mathbf{a}) = 0$. This means that the score of a NE (by either definition) is always strictly better than the score of any profile that is not a NE. Further, any profile that is r -SE has a better strict-stability score than any non r -SE profile. A similar property holds w.r.t. r -SSE. As a different example, a profile \mathbf{a} is *Pareto efficient* in G iff $\mathcal{D}_n(G, \mathbf{a}) = 0$.

3. RESOURCE SELECTION GAMES

In this section we demonstrate how stability scores can be used to measure and compare the stability of different outcomes in a given game. To this end we focus on a very simple parametrized family, where games are known to possess at least one pure equilibrium. A natural choice is the family of *resource selection games* (RSG) with identical resources.

In a RSG there is a set of resources $F = \{1, \dots, m\}$, and a non-decreasing cost function $c : [n] \rightarrow \mathbb{R}_+$, where $[n] = \{1, \dots, n\}$. Each agent $i \in N$ can select exactly one resource j , and suffers a cost (negative utility) of $c(n_j)$, where n_j is the number of agents that selected resource j . RSGs are *potential games* and thus always admit a pure Nash equilibrium. In fact, any NE \mathbf{a} of a RSG $G = \langle F, N, c \rangle$ is a *strong equilibrium* [13], and thus all equilibria have the same (strict) stability score. However, this is no longer true if the games are concatenated in a sequence.

Formally, a *sequential RSG* (SRSG) is a RSG with k steps. Thus a strategy of an agent $a_i \in F^k$ requires selecting one resource in each step (actions may not depend on the previous steps).² We next show that the number of coalitional deviations significantly depends on the played equilibrium. We consider games where $m, n, k \geq 2$, focusing mainly on games with 2 steps.

3.1 Counting deviations: an example

Suppose that $m = 4, n = 6, k = 2$ and that $c(t) = t$ for all $t \leq n$. Any profile in which there are exactly 1 or 2 agents on each resource (in each step) is a Nash equilibrium. However, these equilibria differ in their stability against strict deviation of pairs. Suppose that in the first step agents are partitioned $\{1, 2\}, \{3, 4\}, \{5\}, \{6\}$, and repeat the same actions in the second step. Denote this profile by \mathbf{a} . In this case the pair $\{1, 2\}$ can strictly gain as follows: agent 1 joins agent 5 (or 6) in the first step, and agent 2 joins 5 in the second. Thus the cost for each of the two agents drops from 4 to 3. The pair $\{3, 4\}$ can do the same, thus $\mathcal{SD}_2(G, \mathbf{a}) = 2$.

On the other hand, consider a profile \mathbf{b} where players play in the first step as in \mathbf{a} , and in the second step are partitioned $\{1, 3\}, \{2, 4\}, \{5\}, \{6\}$; then no pair can strictly gain by deviating. Notice though, that this is still not a strong equilibrium, as the coalition $\{1, 2, 3, 4\}$ can still gain (agents 2, 3 deviate in the first step, and 1, 4 in the second), thus $\mathcal{SD}_2(G, \mathbf{b}) = 0$ and $\mathcal{SD}_4(G, \mathbf{b}) = 1$.

Finally, in profile \mathbf{c} agents are partitioned $\{1, 5\}, \{2, 6\}, \{3\}, \{4\}$

²Equivalently, the game can be described as a routing game, with k sequential parts and m parallel edges in each part.

(in the second step), and this is a strong equilibrium, i.e. $\mathcal{SD}_r(G, \mathbf{c}) = 0$ for all r . It therefore follows that w.r.t strict stability scores \mathbf{c} is *more stable* than \mathbf{b} , which is more stable than \mathbf{a} .

Note however that none of these profiles is an SSE or even 2-SSE. More generally, in *any* profile in G there is at least one pair (in fact two) that shares a resource and thus they have a (weak) deviation where just one of them gains. Thus for every profile \mathbf{p} in G , we have that $\mathcal{D}_2(G, \mathbf{p}) \geq 2$.

3.2 Bounding stability scores in two-step RSG

The example above shows that different NE profiles in a particular game may differ in their stability to deviations of pairs or larger coalitions. We want to get a better picture of the gap between the most and least stable NE profiles, focusing on pair deviations. For the results in this section, we will restrict our cost function to be convex.

A nondecreasing cost function $c : [n] \rightarrow \mathbb{R}$ is said to be *convex* if it has an increasing marginal loss; i.e., $c(i+1) - c(i) \leq c(j+1) - c(j)$ for every $i < j$. Note that when facing a convex cost function, agents in an RSG try to minimize the maximal number of agents using a single resource. If the number of agents on every resource is the same, we say that the partition is *balanced*. If these numbers differ by at most one, we say that the partition is *nearly balanced*.

Let G be a two-step game with a convex cost function. Note that when $n \bmod m = 0$, any NE is a balanced partition of agents to resources (in each step). In such partition, no coalition can gain by deviating, as at least one deviating agent will end up paying more in expectation. If, in addition, costs are *strictly* convex, then even weak deviations are impossible. Since in this setting every NE is an SE (and even an SSE), stability scores are trivial. We therefore assume that $n \bmod m = q > 0$.

Let $\hat{\mathbf{a}}$ be the profile with the highest number of pair deviations, and let \mathbf{a}^* be the profile with the lowest number of pair deviations.

$$\text{PROPOSITION 1. } \mathcal{SD}_2(G, \hat{\mathbf{a}}) = \Theta\left(\frac{qn^2}{m^2}\right).$$

PROOF SKETCH OF LOWER BOUND. We note that in $\hat{\mathbf{a}}$ agents play some nearly balanced partition in the first step, and repeat the same partition in the second step. Thus some resources (called *full*) will have $\lceil n/m \rceil$ agents, and the others will have $\lfloor n/m \rfloor$ agents. A crucial observation used in the proof (and in the proofs of the other propositions in this section), is that a pair has a strict deviation if and only if it shares a full resource in both steps. Then (similarly to the example above) one agent switches to a non-full resource in the first step, and the other does the same in the second step. \square

Note that when $q = \Theta(m)$, which is a typical situation, there are over $\Omega\left(\frac{n^2}{m}\right)$ deviating pairs.

We find that the best NE \mathbf{a}^* is significantly better than $\hat{\mathbf{a}}$.

$$\text{PROPOSITION 2. } \mathcal{SD}_2(G, \mathbf{a}^*) = O\left(\frac{n^2}{m^2}\right). \text{ Further, if either } n < m^2 \text{ or } q \leq \frac{m}{2}, \text{ then } \mathcal{SD}_2(G, \mathbf{a}^*) = 0, \text{ i.e. } \mathbf{a}^* \text{ is 2-SE.}$$

In order to achieve the upper bound asserted in the proposition we define a profile that tries to scatter in the second step agents that shared a resource in the first step. As a qualitative conclusion, we see that in order to minimize possible deviations, agents should form a partition in the second step that differs as much as possible from the partition in the first step.

3.3 SRSGs with many steps

The following proposition quantifies the stability score of a random pure NE in a RSG with k steps. Note that the set of pure NEs coincides with the set of profiles that are nearly balanced in each step.

PROPOSITION 3. *Let G be an SRSRG with k steps and a convex cost function, and let \mathbf{a} be a random NE in G . The expected number of deviating pairs in G is $\mathcal{SD}_2(G, \mathbf{a}) \cong \binom{n}{2} (1 - (1 + \alpha)e^{-\alpha})$, where $\alpha = \frac{q(k-1)}{m^2}$.*

We can summarize how the parameters affect stability as follows. If the number of steps k is small, and the number of resources m increases, then $\alpha \rightarrow 0$, and thus $\mathcal{SD}_2(G, \mathbf{a}) \rightarrow 0$ as well (i.e. there are very few pairs that can deviate). Conversely, when the number of steps grows (in particular when $k \gg \frac{m^2}{q}$), then almost every pair can deviate with a high probability.

As a corollary of Proposition 3 when $k = 2$, we get the lower bound of Proposition 1 for the case $q = \Theta(m)$, as

$$\mathcal{SD}_2(G, \hat{\mathbf{a}}) \geq \binom{n}{2} \left(1 - \left(1 - \frac{1}{m}\right) \left(1 + \frac{1}{m}\right)\right) = \Omega\left(\frac{n^2}{m}\right).$$

4. STABILITY SCORES IN AD AUCTIONS

Having showed how stability scores can be used to analyze coalitional stability in simple games, we next turn to prove our main results. We compute the stability scores of the VCG and GSP ad auctions, which are central to the recent literature on economic mechanism design. Since both auctions admit strong equilibria, we do not consider strict deviations, and instead focus our analysis on weak deviations and the scores they induce.

4.1 Ad auctions: model and notations

An ad auction has s slots to allocate, and $n \geq 2s$ bidders,³ each with valuation v_i per click [20]. Every slot $1 \leq j \leq s$ is associated with a click-through rate (CTR) $x_j > 0$, where $x_j \geq x_{j+1}$. For mathematical convenience, we define $x_j = 0$ for every $j > s$. Throughout the paper we make the simplifying assumptions that CTRs are strictly decreasing (i.e., $x_j > x_{j+1}$), and that $v_i \neq v_j$ for all $i \neq j$. We denote by bold letter the corresponding vectors of valuations, CTRs, and bids (e.g. $\mathbf{b} = (b_1, \dots, b_n)$).

A bidder i that has been allocated slot j gains v_i per click (regardless of the slot), and is charged p_j per click. Thus, her total utility is given by $u_i = (v_i - p_j)x_j$.

VCG. In the VCG mechanism every bidder i submits a bid b_i , and the mechanism allocates the j 'th slot, $j = 1, \dots, s$, to the j 'th highest bidder. Each bidder j is charged (per click) for the "harm" she poses to the other bidders, i.e., the difference between the welfare of bidders $k \neq j$ if j is omitted and their welfare when j exists.

It is well known that the VCG mechanism is *truthful*, meaning that reporting true valuations $b_j = v_j$ is a (weakly) dominant strategy for all bidders. In particular, it is a Nash equilibrium.

Suppose that bidders' valuations are sorted in non-increasing order. Assuming truthful bidding (i.e. $b_j = v_j$ for all j), each bidder $i \leq s$ is allocated slot i , and pays

$$p_i^{VCG} = \sum_{s+1 \geq j \geq i+1} \frac{x_{j-1} - x_j}{x_i} \cdot v_j. \quad (1)$$

³When discussing deviating pairs it is sufficient to assume $n > s$, which is a typical situation. Also, all of our results can be easily adjusted to cases with fewer bidders.

GSP. In the GSP auction, slot j is given to the j 'th highest bidder (as in the VCG auction). Denote by j the bidder who is getting slot j . The charge of bidder $j = 1, \dots, s$ equals to the bid of the next bidder; i.e., $p_j = b_{j+1}$. For mathematical convenience, we define $b_{j+1} = 0$ for $j \geq n$.

GSP equilibria. Varian [20] identifies a set of natural Nash equilibria of the GSP auction, termed *envy free NE* or *Symmetric NE* (SNE), which are characterized by a set of recursive inequalities. Varian shows that all SNE's satisfy some very convenient properties. First, in SNE no bidder wants to swap slots with any other bidder.⁴ Second, SNEs are efficient in the sense that bidders with higher valuations always bid higher (and thus get better slots). This allows us to assume that valuations are also sorted in non-decreasing order $v_1 \geq v_2 \geq \dots \geq v_n$. Lastly, SNEs can be easily computed by a recursive formula, which makes them especially attractive for computerized and online settings.

The two equilibria that reside on the boundaries of the SNE set, referred to as *Lower Equilibrium* (LE) and *Upper Equilibrium* (UE), are of particular interest. We denote the LE and UE profiles by $\mathbf{b}^L = (b_i^L)_{i \in N}$ and $\mathbf{b}^U = (b_i^U)_{i \in N}$, respectively. The bids in the LE, for every $2 \leq i \leq s+1$, are given by

$$b_i^L x_{i-1} = v_i(x_{i-1} - x_i) + b_{i+1}^L x_i = \sum_{s+1 \geq j \geq i} v_j(x_{j-1} - x_j).$$

In particular, since CTRs are strictly decreasing, we get that $b_i > b_{i+1}$ for all $i \leq s$. A central result by Varian [20] is that the LE equilibrium induces payments, utilities, and revenue equal to those of the truthful outcome in VCG. It is therefore of great interest to compare the stability of these seemingly identical outcomes in both mechanisms.

The bids in the UE, for every $2 \leq i \leq s+1$, are given by

$$b_i^U x_{i-1} = v_{i-1}(x_{i-1} - x_i) + b_{i+1}^U x_i = \sum_{s+1 \geq j \geq i} v_{j-1}(x_{j-1} - x_j).$$

In the remaining of this section we measure the stability of the VCG and GSP mechanisms. Our results indicate that while the mechanisms have seemingly identical outcomes, for many natural valuation and CTR functions, GSP is far more stable than VCG.

4.2 Deviations in VCG

Recall that the payment for bidder i is a weighted average of reported (and by truthfulness, the actual) values of bidders $i+1 \leq j \leq s+1$ (see Eq. (1)).

We next characterize the structure of a set of deviators R of size r . We say that a coalition R of r bidders has a *potential to deviate* under VCG (or that it is a *potential coalition*), if either: (a) the group R contains exactly r winners (i.e., bidders that are allocated a slot $j \leq s$); or (b) the set R is composed of $t < r$ winners, the first loser, and the $r-t-1$ bidders that directly follow (i.e., bidders $s+1$ through $s+r-t$).

We denote the number of potential coalitions of size r by M_r . We argue that it only makes sense to count potential coalitions when considering a deviation.

To see why, note first that all bidders ranked $s+r$ or worse have no effect on the payment of any other bidder, and can be ignored. Second, the bidders ranked $s+2, \dots, s+r-1$ are only effective if they allow the bidder allocated slot $s+1$ to lower her bid. Thus non-potential coalitions must contain at least one bidder that has no contribution at all to the deviation, and can therefore be ignored.

⁴When swapping with a bidder in a worse slot, this requirement coincides with the one implied by NE. However when swapping with a bidder in a better slot, envy-freeness is slightly stronger.

Note for example that while adding dummy bidders (with valuation 0) increases the total number of coalitions, the number of potential coalitions remains unchanged.

It is easy to verify that there are $\binom{s}{r}$ coalitions of type (a), and $\sum_{t=1}^{r-1} \binom{s}{t}$ coalitions of type (b). Thus $M_r = \sum_{t=1}^r \binom{s}{t}$. Interestingly, in VCG every potential coalition can actually deviate.

PROPOSITION 4. *Under the truthful equilibrium of VCG, denoted by T , any potential coalition has a deviation, i.e., $D_r(VCG, T) = M_r$ for all $2 \leq r \leq s$.*

PROOF. Let R be some potential coalition, and $i^* \in \operatorname{argmin}_{i \in R} v_i$. We call i^* the *indifferent bidder*. Suppose that every agent $i \in R$ reports v'_i so that $v_i > v'_i > v_{i+1}$. Clearly, this has no effect on slot allocation. In coalitions that include only winners, all the agents except agent i^* (which is indifferent) pay strictly less than their original payments, as the payment monotonically depends on the valuations of the other members of R . In potential coalitions other type, where R includes t winners and $r - t$ losers, all t winners strictly gain. \square

4.3 Deviations in GSP

Since LE is a Nash equilibrium, we have that $\mathcal{D}_1(GSP, LE) = SD_1(GSP, LE) = 0$. In fact, as in the VCG mechanism, no coalition has a strict deviation from the LE profile in GSP. This statement is not as trivial in the GSP mechanism, but it follows from Lemma 9 toward the end of this section. The same analysis holds for the UE in GSP. We next turn to evaluate the resistance of GSP to (non-strict) deviations, focusing on the lower equilibrium. As in the previous section, we only count potential coalitions as all other coalitions necessarily contain redundant participants.

Pair deviations: characterization

We begin by characterizing all deviations by pairs of agents.

Lower equilibrium. It is easy to see that for every $i \leq s$, the pair of agents $(i, i + 1)$ (called *neighbors*) can always (weakly) gain as a coalition, by having agent $i + 1$ lowering her bid to b'_{i+1} , so that $b_{i+1} > b'_{i+1} > b_{i+2}$.⁵ In this case, agent $i + 1$ is not affected, but agent i gains the difference $x_i(b_{i+1} - b'_{i+1}) > 0$. It is also clear that bidders ranked $s + 2$ or worse can never be part of a deviating pair. In terms of the stability score, this means that

$$s \leq \mathcal{D}_2(GSP, LE) \leq M_2 = \binom{s+1}{2}.$$

Consider the pair of agents (k, j) , where $k < j \leq s + 1$. We want to derive a sufficient and necessary condition under which the pair (k, j) has a deviation. A simple observation is that given some Nash equilibrium, for an agent i to strictly gain by being allocated a new slot $i' \neq i$, the bid $b_{i'+1}$ must strictly decrease, since otherwise this would also be a deviation for i as a single agent (in contradiction to equilibrium). Therefore, either (1) k moves to a worse slot $k' = j - 1$, and $b'_k < b_k$; or (2) j moves to a better slot $j' = k$, k is pushed down to $k' = k + 1$, and $b'_k < b_k$. However, if j gains in case (2), then this means she is envy in bidder k . This is impossible, as we assumed \mathbf{b} is an SNE. Thus, the only deviation is where $k' = j - 1$; $j' = j$. Further, this is a deviation only if $b_{j-1} > b'_k > b'_j \geq b_{j+1}$. Note that: (i) b'_k can get any value in this range without affecting the utility of k or j , (ii) the utility of j remains the same, and (iii) the most profitable deviation for k is one in which $b'_j = b_{j+1}$ (breaking the tie in favor of j).

⁵The assumption that CTRs are strictly decreasing is required here, as otherwise bidder $i + 1$ may not be able to lower her bid.

The discussion above establishes a necessary condition for a pair deviation, and asserts that in every pair deviation of k, j only agent k can strictly gain, where $k < j$. We next complete the characterization by establishing a sufficient condition for pair deviation.

For the following results, we denote $a = x_{j-1} - x_j$ (for our fixed j), and $w_i = \frac{x_{i-1} - x_i}{x_j}$ for all $i \leq s + 1$.

LEMMA 5. *Suppose that the pair k, j deviates from LE, by moving agent k to slot $k' = j - 1$. Let $u(k), u'(k)$ be the utility of agent k before and after the deviation, then*

$$u(k) - u'(k) \geq \sum_{t=k+1}^{j-1} (x_{t-1} - x_t)(v_k - v_t) - a \cdot v_j + a \sum_{i=j+1}^{s+1} w_i v_i.$$

Moreover, in the optimal deviation for agent k the last inequality holds with an equality.

PROOF. Suppose agent j lowers her bid to $b'_j = b_{j+1} + \epsilon$ where $\epsilon \geq 0$ (so j keeps her slot). For any \mathbf{x}, \mathbf{v} the utility of agent k changes as follows:

$$\begin{aligned} u(k) - u'(k) &= (v_k - b_{k+1})x_k - (v_k - (b_{j+1} + \epsilon))x_{j-1} \\ &= (x_k - x_{j-1})v_k - \sum_{t=k+1}^{s+1} (x_{t-1} - x_t)v_t \\ &\quad + \sum_{i=j+1}^{s+1} \frac{x_{j-1}(x_{i-1} - x_i)}{x_j} v_i + \epsilon x_{j-1} \\ &= \sum_{l=k+1}^{j-1} (x_{l-1} - x_l)v_k - \sum_{t=k+1}^j (x_{t-1} - x_t)v_t \\ &\quad + \left(\frac{x_{j-1}}{x_j} - 1 \right) \sum_{i=j+1}^{s+1} (x_{i-1} - x_i)v_i + \epsilon x_{j-1} \\ &= \sum_{t=k+1}^{j-1} (x_{t-1} - x_t)(v_k - v_t) - (x_{j-1} - x_j)v_j \\ &\quad + \frac{x_{j-1} - x_j}{x_j} \sum_{i=j+1}^{s+1} (x_{i-1} - x_i)v_i + \epsilon x_{j-1} \\ &= \sum_{t=k+1}^{j-1} (x_{t-1} - x_t)(v_k - v_t) - a \cdot v_j + a \sum_{i=j+1}^{s+1} w_i v_i + \epsilon x_{j-1}. \end{aligned}$$

The inequality follows since $\epsilon \geq 0$. In the optimal deviation $\epsilon = 0$ in which case we get an equality. Note that $\sum_{i=j+1}^{s+1} w_i v_i$ is a weighted average of valuations. In particular, it is always between v_{s+1} and v_{j+1} . \square

As a direct corollary from Lemma 5, we get that in LE the pair k, j (where $k < j - 1$), has a deviation *if and only if*

$$\sum_{t=k+1}^{j-1} (x_{t-1} - x_t)(v_k - v_t) < a \cdot v_j - a \sum_{i=j+1}^{s+1} w_i v_i. \quad (2)$$

Upper equilibrium. It is easy to check that a similar characterization to Eq. (2) applies to the UE. However, the conditions differ with respect to bidders that are two positions apart.

PROPOSITION 6. *Given a UE, the pair of agents $i, i + 2$ has a deviation for every $i < s$.*

This result holds under all valuation and CTR functions; hence $\mathcal{D}_2(GSP, UE) \geq 2s - 1$. This means that the UE may be slightly

Valuations	CTR	← concave →	← convex →	
		β -concave	Linear	β -convex
concave	2-concave	All $\binom{s+1}{2}$	All $\binom{s+1}{2}$	-
	Linear	$\Omega(s^2)$	$\Theta(s\sqrt{s})$	$O(s \cdot \log_\beta(s))$
convex	2-convex	-	s	s

Table 1: The table summarizes the number of pairs that have a deviation, i.e., $\mathcal{D}_2(GSP, LE)$. When one function is strictly concave and the other is strictly convex, the score may depend on the exact structure of both functions.

less stable than LE (whose stability is expressed in Theorem 7). Yet, it is not too difficult to show that the number of pair deviations from UE and LE are asymptotically the same. Therefore, in the remainder of this section we focus on stability scores of LE.

Pair deviations: quantification

It turns out that the asymptotic number of pair deviations strongly depends on the shape of both the CTR function and the valuation function. In particular, convexity (as well as concavity and β -convexity) will play a major role in our results. Let g_1, \dots, g_m be a monotonically *nonincreasing* vector.

Similarly to the way defined convex cost functions in Section 3, we say that g is *convex* if it has a decreasing marginal loss; i.e., $g_i - g_{i+1} \geq g_j - g_{j+1}$ for every $i < j$. Similarly, if g has an *increasing* marginal loss then it is *concave*.

Note that linear functions are both convex and concave. A special case of convexity (resp., concavity) is when the marginal loss decreases (resp., increases) exponentially fast.

Let $\beta > 1$. We say that g is β -*convex* if $g_{i-1} - g_i \geq \beta(g_i - g_{i+1})$ for every i . Similarly, g is said to be β -*concave* if $\beta(g_{i-1} - g_i) \leq g_i - g_{i+1}$ for every i .⁶

Intuitively, as either valuations or CTRs are “more” convex,⁷ a bidder who deviates by moving to a lower (i.e., worse) slot faces a more significant drop in her utility. Thus we can hope that pairs that are sufficiently distant from one another will not be able to deviate jointly. This intuition is further formalized and quantified in the remainder of this section. For convenience, the results are summarized in Table 1.

The next proposition demonstrates that convexity induces greater stability.

THEOREM 7. *Suppose that both CTR and valuation functions are convex. The number of pairs with deviations in the Lower equilibrium can be upper bounded as follows.*

- (A) $\mathcal{D}_2(GSP, LE) = O(s\sqrt{s})$.
- (B) if CTRs are β -convex then $\mathcal{D}_2(GSP, LE) = O(s \log_\beta s)$.
- (C) if valuations are β -convex, for any $\beta \geq 2$, then only neighbor pairs can deviate. I.e., $\mathcal{D}_2(GSP, LE) = s$.

We present the proof of the first statement, so as to demonstrate the proof technique.

PROOF OF 7(A). Recall that $a = x_{j-1} - x_j > 0$. A crucial observation is that $\sum_{i=k+1}^{s+1} w_i v_i$ is in fact a weighted average of

⁶Lucier et al. [15] studied GSP auctions with *well-separated* CTR functions, which is a closely related term. In particular, a $\frac{1}{\beta}$ -well separated function is also β -convex.

⁷When referring to convexity of CTR/valuation functions, we only consider the first $s + 1$ values.

valuations, where the weight w_i is proportional to the difference $x_{i-1} - x_i$. Therefore this average is biased toward low values when CTR is convex, and toward high values when it is concave.

Also, since CTRs are convex, we have that for all $i < j$, $x_{i-1} - x_i \geq a$. Thus by Lemma 5,

$$\begin{aligned}
u(k) - u'(k) &\geq a \sum_{t=k+1}^{j-1} (v_k - v_t) - a \cdot v_j + a \sum_{i=j+1}^{s+1} w_i v_i \\
&= a \left(\sum_{t=k+1}^{j-1} (v_k - v_t) + \sum_{i=j+1}^{s+1} w_i v_i - v_j \right) \\
&\geq a \left(\sum_{t=k+1}^{j-1} (v_k - v_t) + \frac{\text{avg}_{s+1 \geq i \geq j+1} (v_i) - v_j}{s+1} \right). \quad (3)
\end{aligned}$$

Therefore, in order to prove that the pair j, k can deviate, it is necessary to show $\sum_{t=k+1}^{j-1} (v_k - v_t) < v_j - \frac{\text{avg}_{s+1 \geq i \geq j+1} v_i}{s+1}$. (4)

We note that under linear CTRs, all inequalities become equalities (in which case Equation (4) is also a sufficient condition). Observe that closer pairs are more likely to deviate. E.g. for pairs s.t. $j = k + 2$, it is sufficient that $v_k - v_{k+1} < v_{k+2} - \frac{\text{avg}_{s \geq t' \geq k+3} v_{t'}}{s}$ to have a

deviation. Let $h = j - 1 - k \geq 1$, and $z = v_k - v_{j-1} = v_k - v_{k+h}$.

From convexity of \mathbf{v} it holds that for all $h' < h$, $\frac{v_k - v_{k+h'}}{h'} \geq \frac{v_k - v_{k+h}}{h} = \frac{z}{h}$, thus for the LHS of Eq. (4),

$$\sum_{t=k+1}^{j-1} (v_k - v_t) \geq \sum_{t=k+1}^{j-1} z \frac{t-k}{h} = \frac{z}{h} \frac{h(h+1)}{2} = \frac{h+1}{2} z. \quad (5)$$

Bounding the RHS of Eq. (4), we have

$$\begin{aligned}
v_j - \frac{\text{avg}_{s+1 \geq i \geq j+1} v_i}{s+1} &\leq v_j - v_{\text{avg}\{s+1 \geq i \geq j+1\}} \quad (\text{convexity of } \mathbf{v}) \\
&\leq v_j - v_{\lceil \frac{j+s}{2} \rceil} = v_j - v_{\lceil j + \frac{s-j}{2} \rceil} \\
&\leq \sum_{i'=1}^{\lceil (s-j)/2h \rceil} (v_{j+(i'-1)h} - v_{j+i'h}) \leq \sum_{i'=1}^{\lceil (s-j)/2h \rceil} (v_k - v_{k+h}), \quad (6)
\end{aligned}$$

which is at most $\lceil \frac{s-j}{2h} \rceil z$. By using the bounds we showed on both sides of the equation, condition (4) implies $h + 1 < \lceil \frac{s-j}{h} \rceil$, which must be false whenever $h + 1 = j - k > \sqrt{s}$. Therefore each winner $k \leq s$ can deviate with at most \sqrt{s} other bidders, and there can be at most $s\sqrt{s}$ such pairs. \square

It is evident from Theorem 7, that convexity can guarantee some level of stability, and further, that “more” convexity can induce more stability. Our next result complements this observation, by showing that *concavity* of valuation and CTR functions affects stability in the opposite direction.

THEOREM 8. *Suppose that both CTR and valuation functions are concave. The number of pairs with deviations in the Lower equilibrium can be lower bounded as follows.*

- (A) $\mathcal{D}_2(GSP, LE) = \Omega(s\sqrt{s})$.
- (B) if CTRs are β -concave for any $\beta > 1$, then $\mathcal{D}_2(GSP, LE) = \Omega(s^2)$ (i.e. a constant fraction of all pairs).
- (C) if valuations are β -concave, for any $\beta \geq 2$, then all pairs can deviate. I.e., $\mathcal{D}_2(GSP, LE) = \binom{s+1}{2} = M_2$.

A linear function is both convex and concave. Therefore, in the special case where both CTRs and valuations are linear, we obtain an asymptotically tight estimation of $\mathcal{D}_2(GSP, LE)$.

Deviations of more than two agents

We first characterize the structure of such deviations.

LEMMA 9. *Suppose that $R \subseteq N$ is a coalition that gains by a deviation, and let b_j, b'_j denote the bids of $j \in R$ before and after the deviation, respectively. Then the following hold:*

- (a) *There is at least one bidder $i^* \in R$ that does not gain anything from the deviation; this bidder is called the indifferent bidder.*
- (b) *There is at least one bidder $f \in R$ s.t. $R \setminus \{f\}$ still has a deviation; this bidder is called a free rider.*
- (c) *For all $j \in R$, either $b'_j < b_j$, or the utilities of all agents in R (including j) are unaffected by the bid of j .*

In order to prove the Lemma, we must show that the bidder that is ranked last among the deviators is an indifferent bidder (i^*). The free rider (f) is either the bidder that is ranked first among the deviators, or some bidder that is isolated of all other deviators. In addition, it is shown that bidders that move to a better slot either strictly lose, or cause some other deviator to strictly lose.

As a direct corollary of Lemma 9, given any coalition R of size ≥ 3 , the coalition $R \setminus \{f\}$ can also deviate. By induction, therefore, a coalition R that can deviate always contains a pair that can deviate. Moreover, by part (c) of Lemma 9, it follows that given a deviating coalition of size ≥ 2 , it can be extended by adding a bidder who does not change her bid. As a result, a set R can deviate if and only if it contains a pair that can deviate. This crucial observation facilitates the computation of the number of deviations by coalitions of size r for any $r \geq 3$.

Recall that M_r denotes the number of potential coalitions of size r , and that under VCG auction all of these coalitions actually have a deviation. Clearly, $\mathcal{D}_r(GSP, LE) \leq M_r$. We next show how the accurate number of coalitions asymptotically depends on the size of the coalition r and on the number of slots s .

PROPOSITION 10. *If both CTRs and valuations are convex, then*

$$\mathcal{D}_r(GSP, LE) \leq M_r \cdot O\left(\frac{r^2}{\sqrt{s}}\right).$$

In contrast, if both CTRs and valuations are concave, then

$$\mathcal{D}_r(GSP, LE) \geq M_r \cdot d \cdot \left(1 - \exp\left(-\Omega\left(\frac{r\sqrt{r}}{\sqrt{s}}\right)\right)\right)$$

for any positive constant $d < 1$.

That is, at least in the convex case the number of potential deviations under GSP is significantly smaller than under VCG.

This result also establishes an almost sharp threshold for the case of linear CTRs and valuations. In particular, for every $r \gg \sqrt[3]{s}$, almost all coalitions of size r can deviate, while the proportion of coalitions of size $r \ll \sqrt[4]{s}$ that can deviate goes to 0 (when r is fixed and as s grows).

Proposition 10 confirms that the GSP auction is far more stable than the VCG auction against collusions of relatively small coalitions (at least when CTR and valuations are convex).

5. ELIMINATING GROUP DEVIATIONS

5.1 VCG with a reserve price

Consider a variant of the VCG mechanism that adds a fixed reserve price c . That is, only bidders that reports a value of c or higher get a slot, and payments are computed ignoring the other bidders (i.e. replacing their values with c). It is easy to verify that truth-telling remains a dominant strategy, and that Proposition 4 remains valid if the values of all bidders are strictly above c . However, a bidder whose value is exactly c will not join any coalition: by lowering her reported value she will lose her current slot for sure, whereas previously she enjoyed a positive utility.

Now, consider a VCG mechanism that chooses a reserve price as follows. With probability q , the reserve price is chosen randomly from a sufficiently large interval, and with probability $1 - q$, it is set to 0. Crucially, the probability distribution of the reserve price is common knowledge, but agents submit their reports before its realization is revealed. Let us denote the proposed mechanism by VCG*. While the proposed adjustment seems small, it results in a dramatic increase of stability.

THEOREM 11. *If $s \geq n$, then truth-telling is a SSE in VCG*.*

PROOF. First observe that VCG* is a lottery over strategyproof mechanisms, thus no agent has an incentive to deviate unilaterally. Suppose by way of contradiction that there exists a deviating coalition, and let R be such a coalition of minimal size. Since R is minimal, the indifferent agent $i^* \in R$ (as defined in Prop. 4) must lower her reported value, otherwise the coalition $R \setminus \{i^*\}$ can also deviate. Assume, therefore, that $v'_{i^*} = v_{i^*} - \epsilon$ for some $\epsilon > 0$. It is easy to verify that i^* cannot gain in any outcome of the mechanism. In contrast, there is a non-zero probability that c is chosen in the range (v'_{i^*}, v_{i^*}) , in which case the utility of i^* becomes 0, compared to $(v_{i^*} - c)x_{i^*} > 0$ under truth-telling. Therefore, agent i^* loses in expectation, contradicting the existence of a coalition R . \square

By the last theorem, VCG* guarantees stability whenever $n \leq s$.⁸ However, if $s < n$ the bidder ranked $s + 1$ can serve as the indifferent bidder of any coalition. Consequently, VCG* does not possess a SSE. That is, since the utility of agent $s + 1$ is always 0, she will not be discouraged by the random reserve price, even when her reported value falls below the reserve price.

In order to deal with the lack of slots (i.e., the case in which $s \leq n$), we introduce a modified VCG* mechanism, which always induces truth-telling as a SSE.

Consider the following modification to VCG*, termed VCG* $_{\lambda}$. Let $0 < \lambda < \frac{1}{n}$. Given some slot $j \leq s$ with a CTR of $x_j > 0$, it is allocated to the bidder that is ranked j with probability $1 - \lambda$, and is allocated to the bidder that is ranked $s + 1$ with probability λ . This modification effectively creates a new slot $s + 1$, whose expected CTR is λx_j , whereas the new (expected) CTR of slot j becomes $(1 - \lambda)x_j$. This procedure can be applied to the desired additional $n - s$ slots. In particular, a possible instantiation is where the new expected CTR of position s will be $(1 - (n - s)\lambda)x_s$, and there will be $n - s$ new slots with an expected CTR of λx_s . Since the new auction has n slots, the mechanism VCG* can be performed to eliminate all coalitional deviations.

The careful reader will notice that by changing the CTRs, the equilibrium in the new auction may change. However, as long as the order of the slots is preserved, the equilibrium allocation is not affected, and this is ensured by satisfying $\lambda < \frac{1}{n}$. Moreover, the new payment differs from the original payment by at most $v_1 \cdot n \cdot \lambda$; thus for a sufficiently small λ the difference is negligible. As a result, we get the following corollary.

⁸The proof in fact shows a stronger result: truth-telling is a SSE in dominant strategies. Thus VCG* is group-strategyproof.

COROLLARY 12. *Truth-telling is a SSE in mechanism VCG_λ^* for every $0 < \lambda < \frac{1}{n}$. Moreover, the payments and revenue of VCG_λ^* can be arbitrarily close to the payments and revenue of VCG.*

5.2 GSP with a reserve price

As evident from the results in the last section, stability of the VCG mechanism is significantly increased by augmenting the mechanism with a random reserve price and additional subtle randomization. It might be tempting to apply the same technique to the GSP mechanism, in an attempt to increase its stability, while maintaining the possibility to achieve a higher revenue than VCG. Unfortunately, this approach fails since (in contrast to VCG) adding a reserve price does not preserve its original set of equilibria.

To see this, consider a GSP mechanism with a fixed reserve price c . Bidder i is affected by the reserve price if either: (I) $v_i > c > b_i$, in which case bidder i has an incentive to raise her bid, as otherwise she will lose the slot; or (II) $v_i < c < b_i$, in which case she has an incentive to lower her bid, as otherwise she will pay more than the slot's worth to her. In both cases it follows that the modified GSP mechanism no longer preserves the SNE properties characterized by Varian (even with respect to unilateral deviations). The reason for the difference between VCG and GSP is that VCG induces truthful revelation in equilibrium; hence cases (I) and (II) suggested above cannot be realized.

6. DISCUSSION AND FUTURE WORK

Our main contribution in this paper is the introduction of *stability scores* — a new stability measure for game equilibria. We demonstrated how stability scores can be used to compare equilibria in congestion games and to draw qualitative results regarding properties of the game and the profiles that increase coalitional stability.

Auctions. Our results indicate that for a prominent class of CTR and valuation functions, GSP is far more stable than VCG.⁹

It is known that the LE of GSP generates exactly the same revenue as VCG, and any other SNE of GSP generates an even higher revenue. This may suggest that GSP is better than VCG with respect to both revenue and stability. However, a relatively simple modification to the VCG mechanism induces a randomized mechanism that eliminates all coalitional deviations, thus turning it into a highly stable mechanism. An open question is whether our results still hold when ads' quality is also considered (see [20]).

Equilibria selection and mechanism design. Analysis of stability scores can be applied to various games and mechanisms. In particular, in games that have multiple Nash equilibria such analysis can aid in selecting the an equilibrium. Understanding how coalitional stability is affected by properties of the game will help us to play better as players, and to create better games as designers.

Toward a realistic picture of coalitional stability. Solution concepts such as ϵ -NE (or ϵ -SE) quantify the benefit an agent or coalition can get from a deviation. Therefore they offer stability under a relaxed notion of self-interest (i.e. agents will only bother to deviate for some substantial gain). In contrast, stability scores still assume purely self-interested agents, but relax a different aspect of coalitional rationality. Practical limitations on information, communication or trust may mean that a coalition of agents will not collude even if they have a potentially high incentive to do so. Other models such as Myerson's [18] assume that limitations on collusion are given in an explicit and structured form.

⁹Empirical studies indicate that CTRs on common platforms are indeed convex, see [5].

In future research we may wish take a combined approach to coalitional stability, considering both known and unknown limitations on collusion, possibly attributing more importance to coalitions with a stronger incentive to deviate. Such models will enable us to better predict realistic outcomes of games, and to improve the mechanisms we design.

7. REFERENCES

- [1] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *Proc. of 18th SODA*, pages 189–198, 2007.
- [2] R. J. Aumann. Acceptable points in general cooperative n-person games. In *Contribution to the Theory of Games, Vol. IV, Annals of Math. Studies, 40*, pages 287–324, 1959.
- [3] Y. Bachrach. Honor among thieves: collusion in multi-unit auctions. In *Proc. of 9th AAMAS*, pages 617–624, 2010.
- [4] B. D. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof Nash equilibria I. Concepts. *Journal of Economic Theory*, 42(1):1–12, 1987.
- [5] N. Brooks. The Atlas rank report: How search engine rank impacts traffic. Available from <http://tiny.cc/wj8f3>.
- [6] V. Conitzer and T. Sandholm. Failures of the VCG mechanism in combinatorial auctions and exchanges. In *Proc. of 5th AAMAS*, pages 521–528, 2006.
- [7] G. Demange. On group stability in hierarchies and networks. *Journal of Political Economy*, 112:754–778, 2004.
- [8] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [9] B. Edelman and M. Schwarz. Optimal auction design and equilibrium selection in sponsored search auctions. *American Economic Review*, 100:597–602, 2010.
- [10] M. Feldman, R. Meir, and M. Tennenholtz. Stability scores: Measuring coalitional stability. *CoRR*, abs/1105.5983, 2011.
- [11] M. Feldman and M. Tennenholtz. Partition equilibrium. In *Proc. of 2nd SAGT*, pages 48–59, 2009.
- [12] D. A. Graham and R. C. Marshall. Collusive bidder behavior at single-object second-price and english auctions. *Journal of Political Economy*, 95:579–599, 1987.
- [13] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21:85–101, 1997.
- [14] D. Kuminov and M. Tennenholtz. User modeling in position auctions: re-considering the GSP and VCG mechanisms. In *Proc. of 8th AAMAS*, pages 273–280, 2009.
- [15] B. Lucier, R. Paes Leme, and É. Tardos. On revenue in the generalized second price auction. In *The 7th ACM EC Workshop on Ad Auctions*, 2011.
- [16] R. P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, 1987.
- [17] R. Meir, J. S. Rosenschein, and E. Malizia. Subsidies, stability, and restricted cooperation in coalitional games. In *Proc. of 22nd IJCAI*, pages 301–306, 2011.
- [18] R. B. Myerson. Graphs and cooperation in games. *Mathematics of operations research*, 2(3):225–229, 1977.
- [19] D. R. M. Thompson and K. Leyton-Brown. Computational analysis of perfect-information position auctions. In *Proc. of 10th ACM-EC*, 2009.
- [20] H. R. Varian. Position auctions. *Int. Journal of Industrial Organization*, 25(6):1163–1178, 2007.