# Incentive Analysis of Approximately Efficient Allocation Algorithms

Yevgeniy Vorobeychik* and Yagil Engel**

**Abstract.** We present a series of results providing evidence that the incentive problem with *VCG*-based mechanisms is not very severe. Our first result uses average-case analysis to show that if an algorithm can solve the allocation problem well for a large proportion of instances, incentives to lie essentially disappear. We next show that even if such incentives exist, a simple enhancement of the mechanism makes it unlikely that any player will find an improving deviation. In the experimental part of the paper, we demonstrate that incentives to lie decrease with increasing problem complexity. However, we also note that if incentives to lie do exist, they can have a negative impact on welfare.

## 1 Introduction

The field of mechanism design has received considerable attention in academic literature in the last several decades. Great technological advances, coupled with a rather mature understanding of the field, have recently brought much of this theory to bear on real allocation problems faced by the government and industry. Perhaps of greatest practical significance has been the field of auction theory [Krishna, 2002], and, in particular, the design of *combinatorial auctions* [Cramton et al., 2006]. In a combinatorial auction, bidders are allowed to submit bids on all subsets of a given set of items. The auctioneer must then solve the *winner determination problem (WDP)*, computing which subsets of the goods will be allocated to which bidders, with the objective of maximizing efficiency. The field of mechanism design has historically occupied itself primarily with the issue of incentives, while mostly ignoring the computational aspects of the problem. As it turns out, computational impediments can be devastating for incentives. For example, while *VCG* [Mas-Colell et al., 1995] is the central mechanism used to incentivize bidders to report their true valuations, using *VCG*-based payment schemes together with an approximate algorithm for the *WDP* nearly universally fails to yield truthful revelation of values [Nisan and Ronen, 2007]. However, it is well known that the combinatorial auction *WDP* is NP-Hard [Lehmann et al., 2006b] —indeed, even hard to approximate [Sandholm, 2002]— and, consequently, an approximation algorithm must, in general, be used.

* University of Pennsylvania, Computer and Information Science Department; yev@seas.upenn.edu

** Technion, Industrial Engineering and Management; yagile@ie.technion.ac.il; supported in part by an Aly Kaufman Fellowship.

For many years, the worst-case complexity of the winner determination problem stood in the way of acceptance of combinatorial auctions in many real domains. However, recent algorithmic advances, as well as the advances in computer technology, have made it quite feasible to run such auctions even on very large scale problems. In practice, it turns out, the majority of problems studied in simulation can be solved very fast using modern algorithms [Leyton-Brown and Shoham, 2006, Sandholm, 2002, Sandholm et al., 2005]; indeed, the general-purpose CPLEX integer programming tool is usually very effective [Sandholm et al., 2005]. Thus, the worst-case complexity of *WDP* seems rarely to manifest itself in practice. Unfortunately, current theoretical results that address bidder incentives only consider worst-case guarantees.

We present a series of results, both theoretical and experimental, that provide evidence that the incentive problem with *VCG*-based mechanisms is not very severe. In particular, if an algorithm can solve the allocation problem exactly in almost every instance, there are no incentives to deviate from truthfulness in the Bayes-Nash sense. Furthermore, we show that the designer can use sampling to reduce the likelihood that any player will compute an improving deviation. Significantly, this is a *typical-case*, and not a *worst-case* result. Finally, we provide experimental evidence that increasing the size of the problem (either the number of bidders or the number of bids) attenuates the incentives of bidders to lie about their valuations.

There has been a rich literature dealing with incentives in combinatorial auctions with approximate allocation algorithms. Nisan and Ronen [2007] develop a second-chance mechanism in which players are not capable of computing a beneficial lie. In many ways, this work is most similar to ours, particularly to our results in Section 6.1. Sanghvi and Parkes [2004] demonstrate that computing an improving deviation in *VCG*-based combinatorial auctions is NP-Hard. Lavi and Swamy [2005] present a truthful (in expectation) mechanism when the approximation algorithm bounds the integrality gap of LP relaxation, while Lehmann et al. [2002] and Mu'alem and Nisan [2002] obtain general truthful mechanisms for combinatorial auctions when bidders are "single-minded" (i.e., each has positive value for exactly one bundle of items). Dobzinski et al. [2006] present a framework for designing truthful approximation algorithms, and demonstrate instances with an asymptotically optimal worst-case bound for the general WDP.

## 2 Preliminaries

Let $O$ be the set of outcomes (e.g., feasible allocations), $I$ be a set of $n$ players and denote by $T_i$ the set of player $i$'s types, with $T = T_1 \times \cdots \times T_n$ the joint type set. Together with a probability distribution over joint player types $F(\cdot)$ and player utility functions $u_i(t_i, o)$ where $o \in O$ depends on joint player actions $a \in A$, this defines a game of incomplete information. Players can condition their actions on their own type realizations, but not those of other players, giving rise to *strategies*, $s_i : T_i \to A_i$, which select an action for any player type. We use $s$ to denote a vector of joint strategies.

A *mechanism* is a tuple $[A, m]$, where $A = A_1 \times \cdots \times A_n$ is the set of actions available to players and $m : A \to O$ is an outcome function which selects some outcome in $O$ for every joint strategic choice $a \in A$ by the players. We focus primarily on *direct revelation mechanisms*, which restrict the strategies of players to be their types, that is, where $A = T$. In our setting players can submit and accept payments, and their utility functions are quasi-linear in these, that is

$$u_i(t_i, o, p_i) = v_i(t_i, o) + p_i,$$

where $v_i$ is the value that player $i$ with type $t_i$ has for outcome $o$, and $p_i$ is his payment.

A central aspect of mechanism design is the prediction of agent play for a given choice of a mechanism. Typically the role of such predictions is played by equilibrium concepts, and we will appeal to two such concepts below (defined with respect to direct revelation mechanisms). Under a *dominant strategy equilibrium* for every player $i$, type $t_i$, and possible report $t'_i$, $u_i(t_i, m(t_i, t_{-i})) \geq u_i(t_i, m(t'_i, t_{-i})) \; \forall t_{-i}$. That is, each player is (weakly) best off reporting their true type *no matter what other players will do*. Under a *Bayes-Nash equilibrium* for every player $i$, type $t_i$, and possible report $t'_i$,

$$E_F[u_i(t_i, m(t_i, t_{-i}))|t_i] \geq E_F[u_i(t_i, m(t'_i, t_{-i}))|t_i].$$

Both equilibrium concepts admit natural notions of approximation: in an $\epsilon$-dominant strategy equilibrium, a player can gain no more than $\epsilon$ by deviating, no matter what the opponents do, whereas an $\epsilon$-Bayes-Nash equilibrium guarantees that the expected gain to any player from deviation is at most $\epsilon$, with expectation taken with respect to the joint type distribution.[1]

A useful measure of strategic stability is that of *game-theoretic regret*. While in general this regret measure can be defined for any joint strategy profile of players, we use it below only to gauge the regret of truthful reporting. Hence, we use a simpler definition, with

$$\tilde{\epsilon} = E_F[\epsilon(t)] = E_F[\max_i \epsilon_i(t_i)],$$

where

$$\epsilon_i(t_i) = \max_{t'_i \in T_i} E_F[u_i(t_i, m(t'_i, t_{-i})) - u_i(t_i, m(t))|t_i].$$

In words, it is the maximum expected benefit any player can obtain from reporting untruthfully.

The social welfare is defined as

$$V(t, o) = \sum_{i \in I} v_i(t_i, o),$$

[1] Our definition of the approximate Bayes-Nash equilibrium does not condition on the types of individual players. Alternatively, one may use a worst-case measure which computes the maximum deviation any player *type* can attain. In the results below we use the former definition.

where $o$ is an outcome and $t$ is a joint type profile. A common goal of mechanism design is to maximize social welfare. Let $o^* : T \to O$ denote the welfare optimal (efficient) outcome (allocation) and let

$$V^*(t) = \sum_{i \in I} v_i(t_i, o^*(t)) = \max_{o \in O} \sum_{i \in I} v_i(t_i, o)$$

be the maximum welfare achieved for a type profile $t$. It is well known that optimal allocation can be achieved as a truthful dominant strategy equilibrium by using Groves payments [Mas-Colell et al., 1995], with

$$p_i(t) = \sum_{j \neq i} v_j(t_j, o^*(t)) + h_i(t_{-i}).$$

Here $h_i$ is some arbitrary function of the types reported by other players.

Let $g : T \times \Omega \to O$ be an algorithm for computing approximately efficient allocation, where $T$ is the player type space, $\Omega$ is a set of random outcomes, and $O$ is the space of feasible allocations. The randomness in this definition does not necessarily imply randomized algorithms (although that's certainly allowed), but may also suggest that the agents are uncertain about the outcome of the algorithm, perhaps because they do not have sufficient time to run it. In any case, it turns out that, assuming risk neutral agents, randomization can be folded into expectations in much of our analysis. We will use $g(t)$ as a shorthand to indicate that expectation of the expression is taken with respect to $\Omega$ and will overload $v_i(\cdot)$ to mean expected valuations. Since $g$ may compute a suboptimal allocation, we let $V_g(t)$ be the welfare at the allocation $g(t)$, that is $V_g(t) = \sum_{i \in I} v_i(t_i, g(t))$. While the *VCG* mechanism above is defined with respect to the optimal welfare, we can apply the idea with approximately optimal functions $g$. We call the resulting mechanism *VCG*-based, borrowing the terminology of Nisan and Ronen [2007]. Define *VCG*-based payments by $p_i^g(t) = \sum_{j \neq i} v_j(t_j, g(t)) + h_i(t_{-i})$. Hence, the *VCG*-based mechanism will select an outcome according to $g$, and the players will receive payments $p_i^g(t)$.

## 3 Welfare Properties of *VCG*-based Mechanisms

It is generally assumed that incentives to lie are undesirable, in part because they result in greater uncertainty about outcomes. However, such incentives would pose a substantially lesser problem if they are aligned with social utility. Note that under *VCG*-based payments any unilateral deviation that improves player $i$'s utility is also welfare improving. Specifically, observe that

$$u_i(t_i, g(t_i', t_{-i})) = \sum_j v_j(t_j, g(t_i', t_{-i})) > u_i(t_i, g(t)) = \sum_j v_j(t_j, g(t)).$$

However, group deviations may in general lead to welfare loss. A key question, then, is whether there necessarily exists a welfare improving Bayes-Nash equilibrium strategy profile. We now show that the answer is, in general, negative.

*Example 1.* Consider the following combinatorial auction setting. We have two players (1 and 2) and two items (1 and 2). As is standard, assume that $v_1(\emptyset) = v_2(\emptyset) = 0$ and consider the following value functions:

$$v_1(\{1,2\}) = 10, \quad v_1(\{1\}) = v_1(\{2\}) = 4$$
$$v_2(\{1,2\}) = 5, \quad v_2(\{1\}) = v_2(\{2\}) = 2.$$

Define $v_1'$ and $v_2'$ to be:

$$v_1'(\{1,2\}) = 2, \quad v_1'(\{1\}) = v_1'(\{2\}) = 0$$
$$v_2'(\{1,2\}) = 2, \quad v_2'(\{1\}) = v_2'(\{2\}) = 0$$

and suppose that the algorithm $g$ allocates the items as follows:

- $g(v_1, v_2)$ assigns good 1 to player 1 and good 2 to player 2 (for a total welfare of 6)
- $g(v_1', v_2)$ assigns both goods to player 1 (to yield the optimal welfare of 10)
- $g(v_1', v_2')$ assigns both goods to player 2 (to yield total welfare of 5)
- $g(v_1'', v_2')$ assigns both goods to player 2 for any $v_1''$
- $g(v_1', v_2'')$ assigns both goods to player 1 for all $v_2''$ except $v_2'' = v_2'$.
- $g(v_1'', v_2'')$ assigns both goods to player 2 for all $v_1''$ and $v_2''$, with the exception of the cases outlined above.

Now, for the computation of player utilities below, ignore the $h_i$ payment term as it does not affect the players' incentives. Observe that in this example, $(v_1, v_2)$ is not an equilibrium, since the utility to each players is 6, whereas player 1 could obtain 10 by deviating to $v_1'$, which would yield the utility of 2 for the second player. Furthermore, $(v_1', v_2)$ is not an equilibrium either, since player 2 could now gain by deviating to $v_2'$, obtaining the utility of 5, which would give player 1 the utility of 2. The profile $(v_1', v_2')$ is, however, an equilibrium, and yields lower welfare than the truthful profile. Additionally, any pure or mixed strategy profile with support on $v_1''$ and $v_2''$ that are not the special cases described above will yield the same welfare as $(v_1', v_2')$.

To see that no mixed strategy equilibrium with any support will do the job, note that we can only increase welfare by having player 2 play $v_2$ as a part of the support. Without loss of generality, let's look at the restricted game with player 1 choosing between actions $v_1$ and $v_1'$ and player 2 choosing between $v_2$ and $v_2'$. Suppose that player 2 plays $v_2$ with probability $\alpha$ and $v_2'$ with probability $1-\alpha$. Then the utility of player 1 from playing $v_1$ is $4\alpha + 2$, while his utility from playing $v_1'$ is $8\alpha + 2$, and the two are only equal when $\alpha = 0$, that is, when player 2 always selects $v_2'$. In this case, all profiles yield welfare of 5. □

We now formally state the negative result demonstrated by the above example.

**Proposition 1.** *Let $P$ be a combinatorial allocation problem. Then there exists an allocation algorithm $g$ and player valuation functions $v_i$ with $v_i(\emptyset) = 0$ for all players $i$ such that every Bayes-Nash equilibrium yields strictly lower welfare than the strategy profile in which all players report their valuations truthfully.*

For the purposes of the above example, we had to construct a rather bizarre outcome function $g(t)$. An open question is whether some typical approximation algorithms have properties which do ensure that at least one equilibrium (or, ideally, all equilibria) is (are) welfare improving as compared to truthful reporting.

## 4 Uniform Approximation Bound

The first step of our endeavor is to assess the incentives to deviate when the algorithm $g$ is an $\alpha$-approximation of the optimal allocation. It turns out that in this case one can bound the incentives of any player to deviate by misreporting his type in a way that is independent of other player's actions. The resulting bound of $\frac{\alpha-1}{\alpha}V^*$ is due to Kothari et al. [2005].

**Theorem 1 ([Kothari et al., 2005]).** *Suppose that the algorithm $g$ is an $\alpha$-approximation. Then truthful reporting is an $\epsilon$-dominant strategy equilibrium for $\epsilon \geq \frac{\alpha-1}{\alpha}V^*$, where $V^* = \sup_{t \in T} V^*(t)$*

The proofs of this and other results are in the appendix.

To see that the bound in Theorem 1 is in general tight for the proposed payment scheme, consider the following example:

*Example 2.* Consider a two-bidder two-item combinatorial auction with

$$v_1(\{1\}) = v_1(\{2\}) = v_1(\{1,2\}) = v_2(\{1\}) = v_2(\{1,2\}) = 1, \quad v_2(\{2\}) = 0.$$

It is known that a greedy algorithm is a 2-approximation for combinatorial auctions with submodular bidders (i.e., each bidder has submodular valuation function) [Lehmann et al., 2006a]. The valuations of the two bidders as described above satisfy the submodularity condition, and the greedy algorithm can in the worst case achieve $\frac{1}{2}$ of optimal welfare. Specifically, allocating item 1 to the second bidder and item 2 to the first results in optimal allocation with value of $V^* = 2$. Greedy, on the other hand, can output the reverse allocation with $V_g = 1$. Suppose that greedy allocations the first item to bidder 1 and second to bidder 2. A best response of bidder 2 would be to report $v_2(\{1\}) = 2$ and $v_2(\{1,2\}) = 2$ instead. In this case, the greedy algorithm will return the optimal allocation and the second bidder will gain utility of 1, which is $\frac{1}{2}V^* = \frac{\alpha-1}{\alpha}V^*$. □

## 5 Non-Uniform Approximation Bound

The theoretical bound above is not very encouraging when the worst-case approximation guarantees are poor. In many interesting settings, worst case approximation bounds for fast algorithms are quite high. For example, a greedy algorithm is a 2-approximation in combinatorial auctions with submodular bidders [Lehmann et al., 2006a]. Many real problems, however, are "easy" in that the optimal or nearly optimal allocation can be found extremely fast in practice. Thus, while an algorithm may prove very bad in the worst case, it may be quite effective in a typical case. Our goal now is to incorporate this "empirical" flavor into the analysis of the incentives to deviate. To do this, suppose that the approximation bound for $g$ is a known function $\alpha(t)$ of joint player types. Thus, for "easy" problems (type profiles), $\alpha(t)$ is low.

Our first result echoes Theorem 1, although we must weaken the approximate equilibrium notion to Bayes-Nash. The reason is that a player $i$ is assumed to only know his own type $t_i$, and will thus only have a distribution over $\alpha(t)$ conditional on $t_i$.

**Theorem 2.** *Suppose that the algorithm $g$ is an $\alpha(t)$-approximation. Then a player $i$ can gain at most $\epsilon_i(t_i)$ when others are playing truthfully, where*

$$\epsilon_i(t_i) = E_{t_{-i}}\left[\frac{\alpha(t)-1}{\alpha(t)}V^*(t)|t_i\right].$$

The proof of this and other results can be found in the appendix of the extended version of this paper.[2]

**Corollary 1.** *Suppose that the algorithm $g$ is an $\alpha(t)$-approximation. Then truthful reporting constitutes an $\epsilon$-Bayes-Nash equilibrium for*

$$\epsilon \geq nE_t\left[\frac{\alpha(t)-1}{\alpha(t)}V^*(t)\right].$$

*Proof.* We measure the expected benefit to deviation using the corresponding game-theoretic regret.

$$E_t[\epsilon(t)] = E_{t_i}[\max_i \epsilon_i(t_i)] \leq \sum_i E_{t_i}[\epsilon_i(t_i)] = nE_{t_i}[\epsilon_i(t_i)] = nE_t\left[\frac{\alpha(t)-1}{\alpha(t)}V^*(t)\right].$$

□

Now, suppose that the space of joint types $T$ can be partitioned into "easy" and "hard" type profiles, that is, $T = \underline{T} \cup \overline{T}$. Let $\underline{\alpha} = \sup_{t\in\underline{T}}\alpha(t)$ and $\overline{\alpha} = \sup_{t\in\overline{T}}\alpha(t)$ and assume that $\underline{\alpha} \leq \overline{\alpha}$. Then

$$\begin{aligned}
E_t\left[\frac{\alpha(t)-1}{\alpha(t)}V^*(t)\right] &= \int_{\underline{T}}\frac{\alpha(t)-1}{\alpha(t)}V^*(t)dF(t) + \int_{\overline{T}}\frac{\alpha(t)-1}{\alpha(t)}V^*(t)dF(t) \\
&\leq \frac{\underline{\alpha}-1}{\underline{\alpha}}\int_{\underline{T}}V^*(t)dF(t) + \frac{\overline{\alpha}-1}{\overline{\alpha}}\int_{\overline{T}}V^*(t)dF(t) \\
&= \frac{\underline{\alpha}-1}{\underline{\alpha}}E_t[V^*(t)] - \frac{\underline{\alpha}-1}{\underline{\alpha}}\int_{\overline{T}}V^*(t)dF(t) + \frac{\overline{\alpha}-1}{\overline{\alpha}}\int_{\overline{T}}V^*(t)dF(t) \\
&\leq \frac{\underline{\alpha}-1}{\underline{\alpha}}E_t[V^*(t)] + \left(\frac{1}{\underline{\alpha}} - \frac{1}{\overline{\alpha}}\right)V^*_{\overline{T}}F(\overline{T}),
\end{aligned}$$

[2] The extended version can be found at `http://www.seas.upenn.edu/~yev/2008/approxvcg.pdf`.

where $V^*_{\overline{T}} = \sup_{t \in \overline{T}} V^*(t)$. Observe that since $\left(\frac{1}{\underline{\alpha}} - \frac{1}{\overline{\alpha}}\right) V^*_{\overline{T}}$ is just a constant (and $V^*(t)$ is bounded on $\overline{T}$), as the probability measure of "hard" instances becomes small, the incentives for players to deviate approach $\frac{\underline{\alpha}-1}{\underline{\alpha}} E_t[V^*(t)]$. In the special case when $\underline{\alpha} = 1$ (that is, easy instances can be solved fast *exactly*) as is the case in many combinatorial auction settings, and when $F(\overline{T}) = 0$, that is, when the probability of drawing a hard problem is 0, truthful reporting is a Bayes-Nash equilibrium. Hence the following corollary.

**Corollary 2.** *Suppose that $\underline{\alpha} = 1$ and $F(\overline{T}) = 0$. Then the strategy $s_i(t_i) = t_i$—that is, truthfully reporting actual preferences—is a Bayes-Nash equilibrium under the allocation algorithm $g$.*

*Proof.* $E_t[\epsilon(t)] = nE_t\left[\frac{\alpha(t)-1}{\alpha(t)} V^*(t)\right] = n\frac{\underline{\alpha}-1}{\underline{\alpha}} E_t[V^*(t)] = 0$. □

An important question that we would like to address is how a mechanism designer would determine $\alpha(t)$ for his choice of approximation algorithm. The simplest approach would be loosely empirical: the designer could rely on empirical studies that demonstrate that the algorithm solves the allocation problem very well (low $\alpha(t)$) for large classes of problems with distributions $F(\cdot)$ most similar to the one his believes to be facing. A naive careful approach may even attempt to estimate $E[\alpha(t)]$ for a distribution $F(\cdot)$ according to which the player types are generated (at least as the designer believes it). The corresponding estimate $\tilde{\alpha}$ would not, however, shed sufficient light on the corresponding regret bound that we have derived here. A small modification, however, would suffice to produce probabilistic bounds on regret based on samples from $F$. We suggest the following approach. Consider the random variable $Z = \frac{\alpha(t)-1}{\alpha(t)} V^*(t)$. Let us generate $K$ sample type profiles $t^j$ according to $F(\cdot)$ and take the sample average of $Z$, $\tilde{Z} = \frac{1}{K}\sum_{i=1}^K Z_i = \frac{1}{K}\sum_{i=1}^K \frac{\alpha(t^i)-1}{\alpha(t^i)} V^*(t^i)$, where $\alpha(t^i)$ is an upper bound on the approximation quality of the allocation algorithm for the given $t^i$ and $V^*(t^i)$ the upper bound on optimal welfare.[3] If we suppose that each player's utility is bounded between 0 and 1, then $0 \leq Z \leq n$ (because $Z$ is bounded by total welfare), and thus, by Hoeffding's inequality

$$\Pr\{E[Z] \leq \tilde{Z} + \epsilon\} \geq 1 - e^{-2K\epsilon^2/n}.$$

By combining with the regret bound from Corollary 1 we obtain the following theorem.

**Theorem 3.** *Suppose that we generate $K$ samples of player types $t$ from $F$ and compute $Z_i = \frac{\alpha(t^i)-1}{\alpha(t^i)} V^*(t^i)$ for each type profile $t^i$ in the sample. Let $\tilde{Z}$ be the sample average over $Z_i$. Then for any $\epsilon > 0$,*

$$E[\epsilon(t)] \leq n(\tilde{Z} + \epsilon)$$

[3] If the optimal allocation is too difficult to compute exactly, we may still derive the corresponding upper bounds. For example, if the allocation problem is an integer program, we can use LP relaxation to obtain an upper bound on optimal welfare.

*with probability at least* $1 - e^{-2K\epsilon^2/n}$.

Alternatively, suppose that we have an $\alpha$-approximation algorithm, with $\alpha$, perhaps, relatively large, but we know that the allocation problem can be solved exactly for a large proportion of the type profiles $t$, and we would like to estimate how large this proportion truly is. Define $\overline{T}$ to be the set of all type profiles $t$ such that $\alpha(t) > 0$, and let $Z = \mathbf{I}\{t \in \overline{T}\}$ (that is, the proportion of sampled types on which the problem is solved optimally). As above, but now with $0 \leq Z \leq 1$, we can obtain by the Hoeffding inequality that

$$\Pr\{E[Z] \leq \tilde{Z} + \epsilon\} \geq 1 - e^{-2K\epsilon^2},$$

where $K$ is the number of samples we take and $\tilde{Z}$ is the proportion of types for which our algorithm did not solve the allocation problem optimally. Combining this with the results described above, we get the following theorem.

**Theorem 4.** *Suppose that we have an $\alpha$-approximation algorithm, and after $K$ samples the proportion of "hard" instances (i.e., those type profiles for which the algorithm computed a suboptimal allocation) was $\tilde{Z}$. Then for any $\epsilon > 0$,*

$$E[\epsilon(t)] \leq n\frac{\alpha - 1}{\alpha}(\tilde{Z} + \epsilon)V^*_{\overline{T}}$$

*with probability at least* $1 - e^{-2K\epsilon^2}$.

## 6 Computing a Better Response

We argued above that truth-telling is an approximate Bayes-Nash equilibrium in many mechanism design settings even when the optimal allocation could not always be computed. This is one piece of positive evidence for the efficacy of *VCG*-based mechanisms. Additionally, Sanghvi and Parkes [2004] show that computing a better response is NP-Hard for a combinatorial allocation problem. To complement this worst-case result, we demonstrate in this section that the designer can, via random sampling of player deviations, ensure that players are unlikely to find improving deviations.

### 6.1 Typical Hardness of a Better Response

In formalizing the notion of "typical" hardness, we begin by following the ideas of Friedgut et al. [2008] and Xia and Conitzer [2008] and presuming that the bidders search for an improving deviation by selecting deviation candidates uniformly randomly from the set of all types. While typical hardness would likely depend on the precise characteristics of both the allocation problem and the approximation algorithm used, we describe now a general technique that the designer can use to make *any* such problem *typically hard* in the sense that a uniform random search by any player would be arbitrarily unlikely to produce an improving deviation vis-a-vis truthfulness. While we acknowledge that restricting player algorithms to uniform random search is very limiting to our analysis, we note that one cannot *in general* do better than by sampling uniformly over the feasible region [Spall, 2003]. Since this argument is still not entirely satisfactory, we extend our results to more general search techniques employed by the players below.[4]

Suppose that players report a type profile $t$ to the designer. Fix some player $i$ and let the designer draw $k$ random instances of $t_i' \in T_i$ uniformly randomly. Let $T_i' \subset T_i$ be the set of $k$ such draws, as well as the reported type $t_i$. Let

$$g_i'(t) = \arg \max_{o=g(t_i',t_{-i})|t_i' \in T_i'} \sum_{j \in I} v_j(t_j, g(t_i', t_{-i})).$$

Finally, repeat this process for every player $i \in I$ and let

$$g'(t) = \arg \max_{i \in I} g_i'(t). \tag{1}$$

Observe that this enhancement can only improve social welfare.

**Theorem 5.** $V(g'(t)) \geq V(g(t))$.

This theorem follows simply from the definition of $g'(t)$.

**Theorem 6.** *Let $g'(t)$ be an allocation mechanism as defined in Equation 1. The probability that some player can find an improving deviation by a uniform random sample from his type set is at most $\frac{n}{k+1}$.*[5]

The restriction of player search to be uniform over the type space is quite significant, and it seems intuitive that sampling a large number of types uniformly should be sufficient to make the problem hard for players more generally. As the following theorem suggests, this result does generalize (in a somewhat weaker form), to a very large class of sampling distributions. Let $u$ denote a utility value which ranges over the real line, and let $G(u)$ be the distribution function of player utilities induced by the *designer*'s search process (e.g., uniform sampling from the type space), whereas $H(u)$ is the distribution function of player $i$'s utilities induced by the player's search.

**Theorem 7.** *Let $U_1 = \{u|G(u) = 1\}$ and suppose that $H(U_1) = 0$. Then* $\lim_{k \to \infty} \int G(u)^k dH(u) = 0$.

The interpretation is that as long as the players do not have a positive probability of reaching a utility that is better than *any* that the designer *can possibly* attain, the designer can use random sampling to effectively eliminate incentives to lie.

[4] One may note a close relationship of our sampling approach to incremental mechanism design [Conitzer and Sandholm, 2007], which searches for possible manipulations and corrects the mechanism to prevent them.

[5] Actually, the distributions used by the players and the designer need not be uniform; the key condition is that they must be identical.

A final result in this section makes a rather intuitive connection between low probability of generating an improving deviation and low expected regret for a player. For simplicity, we assume that there is a uniform bound $\alpha$ on the approximation quality of the algorithm, and thus the most a player can gain by deviating is $\frac{\alpha-1}{\alpha}V^*$. This bound takes no heed of the players' algorithmic capabilities, which we will now bring to bear in refining it.

Suppose that the player is using an algorithm which is either a uniform random search in his space of types, or satisfies the conditions of Theorem 7. This means that the designer can make it arbitrarily unlikely that the player will randomly generate an improving deviation by running his algorithm for a finite number of iterations. Suppose that the probability that the player finds an improving deviation is $\delta \ll 1$. Then, as the theorem below suggests, the player's *effective* regret is a small proportion of his optimal gain from deviation.

**Definition 1.** *Let $\epsilon_i(t'_i, t_i)$ be the gain that player $i$ with type $t_i$ obtains by deviating to $t'_i$. That is, $\epsilon_i(t'_i, t_i) = \max_{t'_i}\{0, u_i(t_i, g(t'_i, t_{-i})) - u_i(t_i, g(t'_i, t_{-i}))\}$. If $F(t'_i)$ is a distribution on $t'_i$ (e.g., as induced by the player's algorithm), then* effective *regret is defined to be $\tilde{\epsilon}_i = E_F[\epsilon_i(t'_i, t_i)]$.*

**Theorem 8.** *Suppose that the player's algorithm finds an improving deviation with probability at most $\delta$. Further, suppose that $g$ is an $\alpha$-approximation of optimal allocation. Then the player's* effective *regret is at most $\frac{\alpha-1}{\alpha}\delta V^*$.*

## 7 Simulations

As suggested by the example in Section 3, when a *VCG*-based mechanism uses an approximate allocation algorithm, there may be some welfare loss due to the misaligned incentives. Specifically, it may be that the equilibrium of the game induced by the *VCG*-based mechanism yields strictly lower welfare as compared to truthful reporting under the same mechanism. In this section we use simulations to assess the amount of welfare loss in a non-truthful equilibrium of the *VCG*-based mechanism, as well as the relationship between incentives to lie and allocation problem complexity. All our simulations are in the context of combinatorial auctions and use a greedy approximation algorithm [Lehmann et al., 2006a].

### 7.1 Welfare Loss

To address the question of welfare loss due to using an approximate allocation algorithm, we generated marginal values for each set of items uniformly randomly, enforcing only the monotonicity constraint on the value function (i.e., $S \subset T \Rightarrow v(S) \leq v(T)$). We performed the following procedure to approximate expected Bayes-Nash equilibrium welfare:

1. Sample a joint type profile from the type distribution
2. Sample a random sequence of joint type *reports* (i.e., action profiles)

3. For each action (report) profile $a$, approximate game-theoretic regret:
   (a) For each player
      i. Generate a sequence of random type deviations
      ii. Select a deviation $t_i^*$ with highest payoff
      iii. Compute the benefit for deviating from action $a_i$ to $t_i^*$
4. Evaluate welfare under the action (report) profile with smallest regret

The approximation of expected equilibrium welfare is then obtained by taking a sample average of welfare values obtained by following this procedure. The idea behind the above algorithm is that we consider a random set of candidate equilibrium profiles and use the one with smallest (approximate) regret as an approximate equilibrium profile. The key simplification of this procedure is that for each sample joint type profile, we are approximating a Nash equilibrium of the resulting *complete information game*, rather than computing optimal deviations for each player with respect to a distribution of competitor types and corresponding strategies. The latter approach would clearly be computationally infeasible, since it would require us to approximate (possibly arbitrary) strategy functions on a high-dimensional space of player valuations.[6]

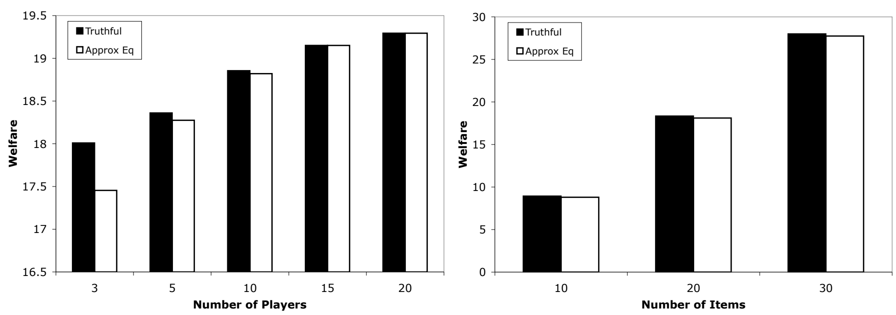


(a) (b)

**Fig. 1.** Average approximate equilibrium welfare in combinatorial auctions with random marginal valuations as a function of (a) players, and (b) items.

Our results, averaged over 100 type profiles for each test point (Figure 1), suggest that equilibrium welfare is, indeed, lower than welfare under truthful reporting. The loss, however, is usually quite small, which is likely due to the fact that expected game-theoretic regret is small for this distribution of valuations, as we observe in the next section. In general these results do suggest that equilibrium welfare can often be lower than truthful, in spite of the fact that unilateral incentives are entirely aligned with social welfare.

[6] Note that the above procedure is in the spirit of an ex-post equilibrium concept. However, we do not enforce a constraint that players cannot condition their strategy on opponents' types.

### 7.2 Incentives to Lie

In this section we use simulations to assess the players' incentives to lie, as well as the relationship of such incentives to the allocation problem complexity. We proceed by generating a random type profile from the joint distribution of player types. Given this fixed profile, we generate random value functions for every player and use these as deviation candidates. We then compute the most gain obtained from 100 such candidates and the final regret value is computed from the fixed type profile by taking the maximum deviation gain obtained by any player.[7] The experiments were repeated for 40-100 different randomly generated type profiles, averaging the regret values over these. The results that we report are actually upper bounds on the incentives to deviate because (a) we consider maximum gain for *every type profile*, rather than for each player given the distribution of other player types, and (b) we report the fraction of utility relative to the greedy, rather than optimal, allocation.

Our first set of simulations considers two kinds of valuations: randomly generated submodular valuations and valuations in which marginal values of items are generated uniformly randomly on the unit interval. We present these results in Figures 2(a) and (b), where we separately vary the number of (a) players and (b) items in the combinatorial auction. As we can observe, in spite of the rather

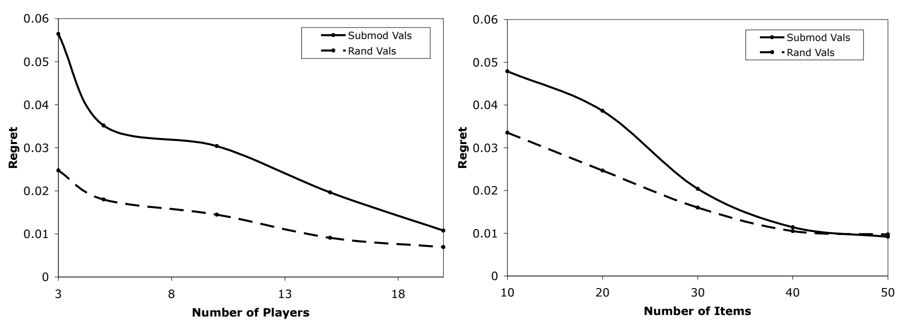


(a) (b)

**Fig. 2.** Average game-theoretic regret, with greedy algorithm approximating WDP of a combinatorial auction with submodular valuations, as a function of the number of (a) players (number of items fixed at 20), (b) items (number of players is fixed at 5).

weak worst-case guarantees on the performance of the greedy algorithm in the case of submodular valuations, the average regret tends to be low. Interestingly, incentives are actually even better when marginal valuations are uniformly random and unconstrained, the case for which no guarantees exist at all for greedy.

[7] It appears that 100 randomly chosen deviations are sufficiently large, as increasing it appears to have insignificant impact on our results. Thus, we are confident that the relationship between problem complexity and regret that we observe below are not a consequence of using a fixed number of deviations.

As we shall see, these results are particular to the value distribution. There is, however, a more robust observation to be made: the incentives to deviate *decrease* with increasing complexity of the allocation problem. In order to assess the robustness of this phenomenon to different distributions of valuations, we use the CATS tool [Leyton-Brown and Shoham, 2006], which generates player valuations for various real-world domains.

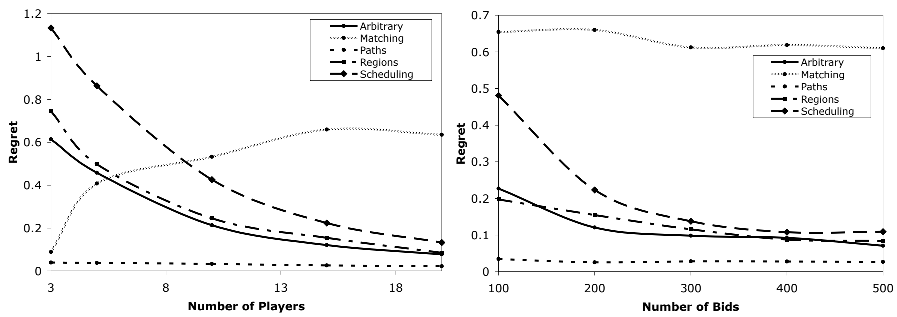


(a) (b)

**Fig. 3.** Average game-theoretic regret in combinatorial auctions with submodular and random valuations, as a function of the number of (a) players (number of items is fixed at 20), (b) bids (number of players is fixed at 15).

The results for several CATS distributions are presented in Figures 3 (a) and (b).[8] Incentives to lie can be relatively large—in fact, in some cases they exceed the total welfare achieved by the greedy algorithm, which is shown to be a very poor heuristics in these domains. However we are interested in the qualitative observation: in all but one case, the incentives to deviate from truthful reporting decrease with increasing problem complexity.

## 8 Conclusion

We presented a series of results that provide evidence that the incentive problem with *VCG*-based mechanisms is not very severe. Our first result used average-case analysis to show that if our algorithm can solve the allocation problem well for a large proportion of instances, incentives to lie essentially disappear. We also showed that even if such incentives exist, the designer can use sampling to make it unlikely that any player will find an improving deviation. This provides a typical-case complement to an already known worst-case hardness result. Finally, we experimentally demonstrate that incentives to lie decrease with increasing problem complexity and show that this result can be observed for a wide range of distributions of combinatorial valuations.

[8] For details about the distributions refer to Leyton-Brown and Shoham [2006].

# Bibliography


Vincent Conitzer and Tuomas Sandholm. Incremental mechanism design. In *Twentieth International Joint Conference on Artificial Intelligence*, pages 1251–1256, 2007.

Peter Cramton, Yoav Shoham, and Richard Steinberg, editors. *Combinatorial Auctions*. MIT Press, 2006.

Shahar Dobzinski, Noam Nisan, and Michael Shapira. Truthful randomized mechanisms for combinatorial auctions. In *Thirty Eighth ACM Symposium on Theory of Computing*, pages 644–652, 2006.

Ehud Friedgut, Gil Kalai, and Noam Nisan. Elections can be manipulated often. In *Forty Ninth Annual Symposium on Foundations of Computer Science*, 2008.

Anshul Kothari, David C. Parkes, and Subhash Suri. Approximately-strategyproof and tractable multi-unit auctions. *Decision Support Systems*, 39:105–121, 2005.

Vijay Krishna. *Auction Theory*. Academic Press, 2002.

Ron Lavi and Chaitanya Swamy. Truthful and near-optimal mechanism design via linear programming. In *Fourty Sixth IEEE Symposium on Foundations of Computer Science*, pages 595–604, 2005.

Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270–296, 2006a.

Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):1–26, 2002.

Daniel Lehmann, Rudolf Müller, and Tuomas Sandholm. The winner determination problem. In Cramton et al. [2006], chapter 12, pages 297–318.

Kevin Leyton-Brown and Yoav Shoham. A test suite for combinatorial auction. In Cramton et al. [2006], chapter 18, pages 451–478.

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *Eighteenth National Conference on Artificial Intelligence*, pages 379–384, 2002.

Noam Nisan and Amir Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.

Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. CABOB: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Science*, 51(3):374–390, 2005.

Saurabh Sanghvi and David C. Parkes. Hard-to-manipulate combinatorial auctions. Technical report, Harvard University, 2004.

James C. Spall. *Introduction to Stochastic Search and Optimization*. Wiley-Interscience, 2003.

Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Ninth ACM Conference on Electronic Commerce*, pages 99–108, 2008.

# A Proofs

## A.1 Proof of Theorem 1

The expected utility of player $i$ in the *VCG*-based mechanism is

$$u_i(t) = \sum_{j \in I} v_j(t_j, g(t)),$$

where we ignore the part of the utility which doesn't depend on $i$'s report $t_i$. Let $t_i^* = \arg\max_{t_i} u_i(t_i, t_{-i})$. The most that the player can gain from deviating to $t_i^*$ is

$$\begin{aligned}
\sum_j v_j(t_j, g(t_i^*, t_{-i})) - \sum_j v_j(t_j, g(t)) \leq \\
&\leq \sum_j v_j(t_j, o^*(t_i^*, t_{-i})) - \sum_j v_j(t_j, g(t)) \\
&\leq \sum_j v_j(t_j, o^*(t_i^*, t_{-i})) - \frac{1}{\alpha} \sum_j v_j(t_j, o^*(t)) \\
&\leq \sum_j v_j(t_j, o^*(t)) - \frac{1}{\alpha} \sum_j v_j(t_j, o^*(t)) \\
&\leq \frac{\alpha - 1}{\alpha} V^*.
\end{aligned}$$

## A.2 Proof of Theorem 2

The proof is quite similar to that for Theorem 1.

Let $t_i^* = \arg\max_{t_i} u_i(t_i, t_{-i})$. The most that the player can gain from deviating to $t_i^*$ is

$$\begin{aligned}
E_{t_{-i}} \left[ \sum_j v_j(t_j, g(t_i^*, t_{-i})) - \sum_j v_j(t_j, g(t)) | t_i \right] \leq \\
&\leq E_{t_{-i}} \left[ \sum_j v_j(t_j, o^*(t_i^*, t_{-i})) - \sum_j v_j(t_j, g(t)) | t_i \right] \\
&\leq E_{t_{-i}} \left[ \sum_j v_j(t_j, o^*(t_i^*, t_{-i})) - \frac{1}{\alpha(t)} \sum_j v_j(t_j, o^*(t)) | t_i \right] \\
&\leq E_{t_{-i}} \left[ \sum_j v_j(t_j, o^*(t)) - \frac{1}{\alpha(t)} \sum_j v_j(t_j, o^*(t)) | t_i \right] \\
&\leq E_{t_{-i}} \left[ \frac{\alpha(t) - 1}{\alpha(t)} V^*(t) | t_i \right].
\end{aligned}$$

## A.3 Proof of Theorem 6

To prove this theorem, we will need the following lemma.

**Lemma 1.** *Let $t_i$ be the type of player $i$ and $t'_i \neq t_i$ be his deviation, and suppose that $t'_i$ is sampled from $T_i$ uniformly randomly. Then for any $q \geq 0$*

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\} = \Pr\{\sum_i v_i(t_i, g(t'_i, t_{-i})) \geq q\}.$$

*Proof.* Observe, first, that for any $q, z$ and any $t'_i \neq t_i$,

$$\Pr\{v_i(t_i, g(t'_i, t_{-i})) \geq q | v_i(t'_i, g(t'_i, t_{-i})) \geq z\} = \Pr\{v_i(t_i, g(t'_i, t_{-i})) \geq q\},$$

since the value function itself is fixed. Consequently, letting $V_{-i}(t'_i) = \sum_{j \neq i} v_j(t_j, g(t'_i, t_{-i}),$ we get

$$\begin{aligned}\Pr\{v_i(t_i, g(t'_i, t_{-i})) + V_{-i}(t'_i) \geq q | v_i(t'_i, g(t'_i, t_{-i})) + V_{-i}(t'_i) \geq z) = \\ E_{V_{-i}(t'_i)=w}[\Pr\{v_i(t_i, g(t'_i, t_{-i})) \geq q - w | v_i(t'_i, g(t'_i, t_{-i})) \geq z - w\}] = \\ E_{V_{-i}(t'_i)=w}[\Pr\{v_i(t_i, g(t'_i, t_{-i})) \geq q - w\}] = \\ \Pr\{v_i(t_i, g(t'_i, t_{-i})) + V_{-i}(t'_i) \geq q\}.\end{aligned}$$

Now, observe that, given the definition of $g'()$,

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\} = \Pr\{\sum_i v_i(t_i, g(t^*_i, t_{-i})) \geq q | v_i(t'_i, g(t^*_i, t_{-i})) + V_{-i}(t^*_i) \geq z\},$$

and, consequently,

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\} = \Pr\{\sum_i v_i(t_i, g(t^*_i, t_{-i})) \geq q\}.$$

But, since both $t'_i$ and $t^*_i$ are drawn uniformly randomly,

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\} = \Pr\{\sum_i v_i(t_i, g(t'_i, t_{-i})) \geq q\}.$$

□

We are now ready to prove the theorem. First, consider a select player $i$.

$$\begin{aligned}\Pr\{u_i(t_i, g'(t'_i, t_{-i})) > u_i(t_i, g'(t))\} &\leq \Pr\{u_i(t_i, g'(t'_i, t_{-i})) \geq u_i(t_i, g'(t))\} \\ &= \Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq \sum_i v_i(t_i, g'(t))\}.\end{aligned}$$

First, consider the quantity $\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\}$ for some fixed $q$. By Lemma 1,

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq q\} = \Pr\{\sum_i v_i(t_i, g(t'_i, t_{-i})) \geq q\}.$$

Hence,

$$\Pr\{\sum_i v_i(t_i, g'(t'_i, t_{-i})) \geq \sum_i v_i(t_i, g'(t))\} = \Pr\{\sum_i v_i(t_i, g(t'_i, t_{-i})) \geq \sum_i v_i(t_i, g'(t))\}.$$

Let $V'(t) = \sum_i v_i(t_i, g(t'_i, t_{-i}))$. By conditioning, we have

$$
\begin{aligned}
\Pr\{\sum_i v_i(t_i, g(t'_i, t_{-i})) \geq \sum_i v_i(t_i, g'(t))\} &= E_{V'(t)}[\Pr\{\sum_i v_i(t_i, g'(t)) \leq z | V'(t) = z\}] \\
&\leq E_{V'(t)}[\Pr\{\max_{t''_i \in T_i \setminus t_i} \sum_i v_i(t_i, g(t''_i, t_{-i})) \leq z | V'(t) = z\}] \\
&= E_{V'(t)}[\Pr\{\sum_i v_i(t_i, g(t''_i, t_{-i})) \leq z | V'(t) = z\}^k].
\end{aligned}
$$

Letting $F(z)$ be the uniform distribution (or any distribution, so long as it is identical for $t'_i$ and $t''_i$), we have

$$
E_{V'(t)}[\Pr\{\sum_i v_i(t_i, g(t''_i, t_{-i})) \leq z | V'(t) = z\}^k] = \int F(z)^k dF(z) = \frac{1}{k+1}.
$$

To conclude the proof, we only need to apply the union bound to obtain $\frac{n}{k+1}$ as the upper bound over all the players.

## A.4 Proof of Theorem 7

First, observe that by the Dominated Convergence Theorem,

$$
\lim_{k \to \infty} \int G(u)^k dF(u) = \int \lim_{k \to \infty} G(u)^k dF(u),
$$

since both $G(u)$ and $F(u)$ are probability densities. Now,

$$
\begin{aligned}
\int \lim_{k \to \infty} G(u)^k dF(u) &= \int_{\overline{U_1}} \lim_{k \to \infty} G(u)^k dF(u) + \int_{U_1} \lim_{k \to \infty} G(u)^k dF(u) \\
&= \int_{\overline{U_1}} \lim_{k \to \infty} G(u)^k dF(u).
\end{aligned}
$$

Since for every $u \in \overline{U_1}$, $\lim_{k \to \infty} G(u)^k = 0$,

$$
\int \lim_{k \to \infty} G(u)^k dF(u) = \int_{\overline{U_1}} \lim_{k \to \infty} G(u)^k dF(u) = 0.
$$

## A.5 Proof of Theorem 8

Let $T_i = \underline{T}_i \cup \overline{T}_i$, with $\overline{T}_i = \{t'_i \in T_i | u_i(t_i, g(t'_i, t_{-i}) > u_i(t_i, g(t))\}$ and $\underline{T}_i = T_i - \overline{T}_i$. Hence, if $F$ is the probability distribution of $t'_i$, $F(\overline{T}_i) \leq \delta$.

By definition of effective regret,

$$
\begin{aligned}
\tilde{\epsilon}_i &= \int_{T_i} \epsilon_i(t'_i, t_i) dF(t'_i) = \int_{\overline{T}_i} \epsilon_i(t'_i, t_i) dF(t'_i) + \int_{\underline{T}_i} \epsilon_i(t'_i, t_i) dF(t'_i) \\
&= \int_{\overline{T}_i} \epsilon_i(t'_i, t_i) dF(t'_i) \leq \frac{\alpha - 1}{\alpha} V^* \int_{\overline{T}_i} dF(t'_i) \leq \frac{\alpha - 1}{\alpha} \delta V^*,
\end{aligned}
$$

where the third line is due to the fact that $\epsilon_i(t'_i, t_i) = 0$ whenever $t'_i$ is not an improving deviation.