Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions

Adrian Vetta*

November 20, 2003

Abstract. We consider the following class of problems. The value of an outcome to a society is measured via a submodular utility function (submodularity has a natural economic interpretation: decreasing marginal utility). Decisions, however are controlled by non-cooperative agents who seek to maximise their own private utility. We present, under some basic assumptions, guarantees on the social performance of Nash equilibria. For submodular utility functions, any Nash equilibrium gives an expected social utility within a factor 2 of optimal, subject to a function-dependent additive term. For non-decreasing, submodular utility functions, any Nash equilibrium gives an expected social utility within a factor 2 of optimal. A condition under which all sets of social and private utility functions induce pure strategy Nash equilibria is presented. The case in which agents, themselves, make use of approximation algorithms in decision making is discussed and performance guarantees given. Finally we present some specific problems that fall into our framework. These include the competitive versions of the facility location problem and k-median problem, a maximisation version of the traffic routing problem of Roughgarden and Tardos [16], and multiple-item auctions.

1. Introduction

Computer scientists have long studied the costs incurred by the lack of *complete information* or the lack of *unbounded computational resources*. For example, the fields of on-line algorithms and approximation algorithms were developed in response to these two problems. However, these fields presume a single authority (or agent) whose goal is to optimise some objective function. What happens when there is a clear social objective function but no single authority? In particular, what if there are many agents whose goals are to optimise their own private objective functions, rather than to collectively optimise the social objective function? Motivated by examples of this type concerning the internet, Koutsoupias and Papadimitriou [8] proposed applying game-theoretic techniques in order to analyse the costs resulting from a lack of *coordination*. Specifically, they proposed the study of non-cooperative games via the use of Nash equilibria (where the agents' strategies are mutual best responses to each other). Given the non-cooperative nature of these

*School of Computer Science and Department of Mathematics, McGill University. Email: `vetta@math.mcgill.ca`

games and the fact that such games may have many Nash equilibria, they proposed studying such equilibria from a worst case perspective. That is, how bad can a Nash equilibrium be, with respect to the social objective, in comparison to the best cooperative solution (or solution produced in presence of a single authority).

The study of Nash equilibria is especially fruitful for problems in which the actions of the agents may be changed quickly and at little cost. This is because it is in such circumstances that Nash equilibria are most likely to arise in practice. Such problems abound in the high-tech economy. From a theoretical viewpoint, notable amongst them is the traffic routing problem which has been studied with great success by Roughgarden and Tardos [16].

In this paper we consider a large class of problems with the following structure. Decisions are made by a set of non-cooperative agents whose action spaces are subsets of an underlying groundset. The actions of the agents induce some social utility, measured by a set function. The goal of the agents, though, is not to maximise the overall social utility; rather, they seek to maximise their own private utility functions. The only assumptions we make are

- The social utility and private utility functions are measured in the same standard unit.

  This standard utility unit may be money, gold, cake etc. Clearly, such a condition is necessary. For example, no guarantees can be obtained if the value to society is measured in terms of the number of oranges but the agents seek to maximise the number of apples.

- The social utility function is submodular.

  Submodularity corresponds to a property that arises frequently in economics: *decreasing marginal utility.* Here, the additional value accruing from an action decreases as the overall level of activity in the society rises. For example, the additional benefit to a town of an extra taxi company is greater if there are currently no taxi firms in the town rather than if there are already one hundred taxi firms.

- The private utility of an agent is at least the change in social utility that would occur if the agent declined to participate in the game.

  We remark that, equivalently, we require that the private utility of an agent is at least the *Vickrey utility* with respect to that agent. This concept is often considered in the study of auction mechanisms (see Vickrey-Clarke-Groves payment mechanisms). Moreover, this condition arises in other practical situations as we will see in our examples. (To illustrate why this condition is often satisfied, suppose that the members of the society are somehow able to negotiate how the total social utility is divided up amongst the members. Now, given an outcome, consider the change in the value of the game that occurs if agent i then declines to participate in the game. Observe that the other agents will be willing to pay agent i up to this amount just to participate in the game. Thus, this will be the minimum payoff that agent i will accept.)

Problems for which these three assumptions hold are called *utility systems*. For a utility system, it is possible to provide some strong guarantees concerning the social utility provided by any Nash equilibrium (we will also show that good guarantees arise if we relax the third assumption). Specifically, for non-decreasing, submodular objective functions, any Nash equilibrium will give a solution with expected social utility within a factor 2 of the optimal solution, Hence, any Nash equilibrium is always at least half as good as the optimal social solution. For submodular functions in general, the expected social utility of a Nash equilibrium is within a factor 2 of optimal, subject to a function-based additive term (which, as we will see in our examples, often has a clear economic interpretation). An alternative form of guarantee that has interesting interpretations in certain problems (for example, the traffic routing problem) is also given. These results are shown to be tight.

The other main result in the paper is to show that, given a simple condition, utility systems have the desirable property that they possess pure strategy Nash equilibria. We also discuss and provide performance guarantees for instances in which the agents apply approximation algorithms in determining their strategies.

An outline of the paper is as follows. In Section 2 we introduce the necessary background on game theory and submodular functions, and give a toy example to illustrate our ideas. In Section 3 we prove our results concerning the social performance of Nash equilibria. In Section 4 we discuss pure strategy Nash equilibria and mixed strategy Nash equilibria. We then present the simple condition under which a utility system will have pure strategy Nash equilibria. In Section 5 we relax our third assumption and present results for the situation in which the private utility of an agent is comparable to the Vickrey utility with respect to that agent (loss in social utility that would result from the agent dropping out of the game). Since our three assumptions concerning the utility system are not very restrictive, the results are widely applicable. We illustrate this by presenting a range of problems that fit into our framework. Our first examples are competitive versions of the facility location problem and the k-median problem, which we introduce in Section 6. One implication of the results in this section is that competitive markets are less efficient in industries with high fixed costs and high marginal profits. Practical examples of such social inefficiencies include the duplication of work, as well as the over-supply of lucrative markets (and under-supply of less valuable markets) by firms. Our next example, given in Section 7, concerns traffic routing in networks. In Section 8, we consider the issues of polynomial time implementations. These issues include the time it takes to obtain Nash equilibria and also the consequences of agents using approximation algorithms for strategy determination. One example in which speed considerations are of great importance is auctions. Thus, our last example, given in Section 9, is that of multiple-item auctions. We present a simple polynomial time auction that fits into our overall framework. It follows that the allocation of items given by the auction (in the presence of competing agents who bid in a greedy manner) is at least half as efficient as the optimal allocation given by a single authority. This matches the performance guarantee Lehmann, Lehmann and Nisan [9] gave for the problem where a single authority chooses an allocation.

2. Background and a Simple Example

In this section we present the required concepts and terminology. We will illustrate these concepts using the following simple *stable marriage* game. We have a group of men and a group of women which act as a vertex set in a bipartite graph. There is an edge between man i and women j whose value represents the "quality" of resultant relationship should i and j decide to marry. We remark that the bipartite graph need not be complete. We may assume that we have returned to the 1950s and societal norms only allow men to propose to women. The objective of each man and woman is to maximise the value of any marriage. Thus, on receiving a set of propositions, each women will accept the proposal of highest value.

2.1. some game theory.

Suppose we have k agents and disjoint groundsets $V_1, V_2, \ldots, V_k$. Each element in V_i represents an act that agent i may make, $1 \leq i \leq k$; let $a_i \subseteq V_i$ be an *action* (set of acts) available to agent i. We may wish to restrict the set of actions an agent may make; thus we may not allow every subset of V_i to be a feasible action. Towards this end, we let $\mathcal{A}_i = \{a_i \subseteq V_i : a_i \text{ is a feasible action}\} = \{a_i^1, a_i^2, \ldots, a_i^{n_i}\}$ be the set of all actions available to agent i. We call $\mathcal{A}_i$ the *action space* for agent i. In our marriage game the k agents are the men, and the groundset V_i is the set of edges $\delta(i)$ incident to man i in the graph. An action for man i is either just a choice of edge $(i,j) \in \delta(i)$ (corresponding to a proposal to woman j) or the null-choice (making no proposal). Thus $A_i = \{\emptyset\} \cup \{(i,j) \in \delta(i)$. Observe that we could time-warp back further and allow men to propose to and marry multiple women (although each woman may still only accept one proposal); in such a circumstance the action set A_i would consist of all subsets of $\delta(i)$. In our game multiple proposals are not allowed, but this observation is important because it allows us to evaluate non-feasible outcomes, and this will be important for the class of valuation functions that we consider.

A *pure strategy* is one in which the agent decides to carry out a specific action. [For example, in the marriage game a pure strategy for man i corresponds to making a proposal to woman j.] A *mixed strategy* is one in which the agent decides upon an action according to some probability distribution. The *strategy space* $\mathcal{S}_i$ of agent i is the set of mixed strategies. Thus we may represent $\mathcal{S}_i$ as

$$\mathcal{S}_i = \{s_i \in \mathbb{R}^{n_i} : \sum_{j=1}^{n_i} s_i^j = 1, s_i^j \geq 0\}$$

Thus $s_i \in \mathcal{S}_i$ corresponds to the mixed strategy in which action a_i^1 is implemented with probability s_i^1, action a_i^2 is implemented with probability s_i^2, etc. Hence, a pure strategy corresponds to $(0,1)$-vector in $\mathcal{S}_i$. Now let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_k$ and let $\mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2 \times \cdots \times \mathcal{S}_k$. In addition, we let $V = V_1 \cup V_2 \cup \cdots \cup V_k$. Then for a function $f : 2^V \to \mathbb{R}$, we define $\bar{f} : \mathcal{S} \to \mathbb{R}$ as follows

$$\bar{f}(S) = \sum_{A \in \mathcal{A}} f(A) \Pr(A|S)$$

where $\Pr(A|S)$ is the probability that action set $A = \{a_1, a_2, \ldots, a_k\}$ is implemented given that the agents are using the strategy set $S = \{s_1, s_2, \ldots, s_k\}$. Thus $\bar{f}(S)$ is just the expected value of f on the strategy set S.

Given an action set $A = \{a_1, a_2, \ldots, a_k\} \in \mathcal{A}$, let $A \oplus a_i'$ denote the action set obtained if agent i changes its action from a_i to a_i'. Formally, $A \oplus a_i' = \{a_1, \ldots, a_{i-1}, a_i', a_{i+1}, \ldots, a_k\}$. Similarly, given a strategy set $S = \{s_1, s_2, \ldots, s_k\} \in \mathcal{S}$, let $S \oplus s_i' = \{s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_k\}$, i.e. the strategy set obtained if agent i changes its strategy from s_i to s_i'.

In this paper we will denote by $\gamma : 2^V \to \mathbb{R}$ the social utility function. In addition, for each agent $1 \leq i \leq k$, there is a private utility function $\alpha_i : 2^V \to \mathbb{R}$. [For the marriage game the social utility function will be the sum of the the values of each marriage. The private utility of each agent (man) will be the value of the marriage he is in (or sum of value of his marriages in the multiple proposal version).] The goal of each agent is, therefore, to select a strategy in order to maximise its private utility. Clearly, though, such strategies may not produce a good solution with respect to social utility γ. We say that set of strategies $S \in \mathcal{S}$ is a *Nash equilibrium* if no agent has an incentive to change strategy. That is, for any agent i,

$$\bar{\alpha}_i(S) \;\geq\; \bar{\alpha}_i(S \oplus s_i') \qquad \forall s_i' \in \mathcal{S}_i$$

Equivalently, given the other agents strategies, s_i is the best response of agent i. We say that a Nash equilibrium $\{s_1, s_2, \ldots, s_k\}$ is a *pure strategy Nash equilibrium* if, for each agent i, s_i is a pure strategy. Otherwise we say that the Nash equilibrium is a *mixed strategy Nash equilibrium.* The following result is due to Nash [10].

Theorem 2.1. *Any finite, k-person, non-cooperative game has at least one Nash equilibrium.* □

Therefore, the task of comparing the performance of Nash equilibria against a socially optimal solution is feasible.

2.2. SUBMODULAR FUNCTIONS.

A function with the form $f : 2^V \to \mathbb{R}$ is called a *set function.* A set function f is *submodular* if

$$f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y) \qquad \forall X, Y \subseteq V$$

It is *supermodular* if this inequality is reversed. A set function f is *non-decreasing* if $f(X) \leq f(Y)$, $\forall X \subseteq Y \subseteq V$. For a set function f, the *discrete derivative* at $X \subseteq V$ in the direction $D \subseteq V - X$ is defined as

$$f'_D(X) = f(X \cup D) - f(X)$$

The following result is standard. Condition (III) shows that, in economic terms, submodularity corresponds to the property of decreasing marginal utility, that is, the additional value accruing from an action decreases as the overall level of activity in the society rises.

Lemma 2.2. *The following are equivalent:*

(I) *f is submodular.*
(II) *$A \subseteq B$ implies $f'_D(A) \geq f'_D(B)$, $\forall D \subseteq V - B$.*
(III) *$A \subseteq B$ implies $f'_i(A) \geq f'_i(B)$, $\forall i \in V - B$.*

Proof. First, we show that (III) implies (II). So assume that $f'_i(A) \geq f'_i(B)$ when $A \subseteq B$ and $i \in V - B$. Let $D = \{i_1, i_2, \ldots, i_r\} \subseteq V - B$. Then

$$\begin{aligned}
f(A \cup \{i_1\}) - f(A) &\geq f(B \cup \{i_1\}) - f(B) \\
f(A \cup \{i_1, i_2\}) - f(A \cup \{i_1\}) &\geq f(B \cup \{i_1, i_2\}) - f(B \cup \{i_1\}) \\
&\vdots \\
f(A \cup \{i_1, \ldots, i_r\}) - f(A \cup \{i_1, \ldots, i_{r-1}\}) &\geq f(B \cup \{i_1, \ldots, i_r\}) - f(B \cup \{i_1, \ldots, i_{r-1}\})
\end{aligned}$$

Summing, we obtain

$$\begin{aligned}
f(A \cup D) - f(A) &\geq f(B \cup D) - f(B) \\
f'_D(A) &\geq f'_D(B)
\end{aligned}$$

Next, we show that (II) implies (I). Assume that $f'_D(A) \geq f'_D(B)$ when $A \subseteq B$ and $D \subseteq V - B$. Take sets X and Y. Set $A = X \cap Y$, $B = Y$ and $D = X - Y$. Then

$$\begin{aligned} f'_D(A) &\geq f'_D(B) \\ f(D \cup A) - f(A) &\geq f(D \cup B) - f(B) \\ f(X) - f(X \cap Y) &\geq f(X \cup Y) - f(Y) \end{aligned}$$

Finally, we show that (I) implies (III). Assume f is submodular. Thus, $f(X) + f(Y) \geq f(X \cap Y) + f(X \cup Y)$, $\forall X, Y \subseteq V$. Let $A \subseteq B$ and take $i \in V - B$. Set $X = A \cup \{i\}$ and $Y = B$. Then $f(A \cup \{i\}) + f(B) \geq f(A) + f(B \cup \{i\})$. Hence, $f(A \cup \{i\}) - f(A) \geq f(B \cup \{i\}) - f(B)$. □

Observe that the objective function in the marriage game is submodular. To see this consider two sets of proposals P_1 and P_2, where the edge set corresponding to P_1 is a subset of that corresponding to P_2. Now it is easy to see that the increase in the total value of marriages resulting from adding an additional proposal (edge) to P_1 and P_2 is greater (or equal) in the former case.

2.3. UTILITY SYSTEMS.

Given our competitive game, let the optimal social solution be $\Omega = \{\sigma_1, \sigma_2, \ldots, \sigma_k\}$, with optimal value $\text{OPT} = \gamma(\Omega)$. Here we consider the private utilities of the agents in a solution $S = \{s_1, s_2, \ldots, s_k\}$. First, we introduce some more notation. We denote by $\emptyset_i$, the *null strategy (action)* for agent i; such a strategy corresponds to agent i declining to take part in the game. We denote by $\emptyset = \{\emptyset_1, \emptyset_2, \ldots, \emptyset_k\}$ the strategy set in which each player has a null strategy. For simplicity, we will assume that $\gamma(\emptyset) = 0$.

Now take an arbitrary ordering of the agents. Without loss of generality, we may assume that the ordering is $\{1, 2, \ldots, k\}$. Now given $A \in \mathcal{A}$ we set $A^i = \{a_1, a_2, \ldots, a_i, \emptyset_{i+1}, \ldots, \emptyset_k\}$. Similarly given $S \in \mathcal{S}$ we set $S^i = \{s_1, s_2, \ldots, s_i, \emptyset_{i+1}, \ldots, \emptyset_k\}$. Then, by construction, it follows that

Lemma 2.3. *For an action set $A \in \mathcal{A}$ and set function γ, we have $\gamma(A) = \sum_{i=1}^{k} \gamma'_{a_i}(A^{i-1})$.* □

Corollary 2.4. *For any strategy set $S \in \mathcal{S}$ and set function γ, we have $\bar{\gamma}(S) = \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S^{i-1})$.* □

Now take our submodular, social utility function $\gamma : 2^V \to \mathbb{R}$ (we remark that $\bar{\gamma}$ is also submodular) and our collection of private utility functions $\alpha_i : 2^V \to \mathbb{R}$, $1 \leq i \leq k$. Recall that our third assumption regarding the utility functions states that the private utility to an agent is at least as great as the loss in social utility resulting from the agent dropping out of the game. That is, the system $(\gamma, \cup_i \alpha_i)$ has the property

$$\bar{\alpha}_i(S) \geq \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \qquad \forall S \in \mathcal{S} \tag{1}$$

Given condition (1), we say that the system $(\gamma, \cup_i \alpha_i)$ is a *utility system.* Note that for the marriage game this condition holds, if a player drops out of the game the overall value of the game falls by at most the value of his marriage (in fact, possibly less as then his wife may be able to marry another man). The utility system $(\gamma, \cup_i \alpha_i)$ is said to be *basic* if we have equality in condition (1), that is $\bar{\alpha}_i(S) = \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)$. Observe that, since we are assuming that utilities are measured in the same units, we may view the game in the following manner. The function γ represents the value of the game (or size of the cake), and α_i represents the return to the agent i (i.e. the size of agent i's piece of the cake). Therefore we also require that the sum of the sizes of the pieces must be smaller than the total size of the cake. That is we require that the sum of the private utilities of the agents is at most the social utility

$$\sum_i \bar{\alpha}_i(S) \leq \bar{\gamma}(S) \qquad \forall S \in \mathcal{S} \tag{2}$$

In such a circumstance we say that the utility system $(\gamma, \cup_i \alpha_i)$ is *valid.* Note that the utility system for the marriage game is valid. Note, we do not require that $\sum_i \bar{\alpha}_i(S) = \bar{\gamma}(S)$. In fact, as we shall see the value $\bar{\gamma}(S) - \sum_i \bar{\alpha}_i(S)$ often has a clear meaning. For the moment we may view $\bar{\gamma}(S) - \sum_i \bar{\alpha}_i(S)$ as the utility of some non-agent, say the utility of the general public.

Observe that conditions (1) and (2) must hold if the following two conditions hold

$$\alpha_i(A) \geq \gamma'_{a_i}(A \oplus \emptyset_i) \tag{3}$$

$$\sum_i \alpha_i(A) \leq \gamma(A) \tag{4}$$

We now show that valid utility systems do exist.

Theorem 2.5. *For any submodular function γ, there exist functions α_i, $1 \leq i \leq k$ such that $(\gamma, \cup_i \alpha_i)$ is a valid utility system. In particular, the basic utility system is valid.*

Proof. So we need to show that, for the basic utility system, Condition (2) holds. Now

$$\begin{aligned}
\bar{\gamma}(S) &= \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S^{i-1}) \quad \text{[by Lemma 2.3]} \\
&= \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S^{i} \oplus \emptyset_i) \\
&\geq \sum_{i=1}^{k} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \quad \text{[by Lemma 2.2]} \\
&= \sum_{i=1}^{k} \bar{\alpha}_i(S) \quad \text{[since } (\gamma, \cup_i \alpha_i) \text{ is basic]}
\end{aligned}$$

□

3. Main Results

In this section we present our guarantees concerning the social value of a Nash equilibrium. In particular, for a valid utility system with a non-decreasing, submodular, social utility function we will show that any Nash equilibrium has an expected social value of at least half that of an optimal social solution. In fact, following an approach of Conforti and Cornuéjols [2], we obtain a tighter bound (although it provides the same guarantee in the worst case) with respect to a parameter based upon the discrete curvature of the non-decreasing, submodular function. For a valid utility system with a submodular, social utility function it is not possible to obtain a simple multiplicative guarantee. However, the expected social value of the Nash equilibrium is at least half the social optimal subject to an additive term. This additive term is function-dependent and often has a clean social/economic interpretation; for example, we will see in the Section 6 that, for the competitive facility location problem, it is bounded by the fixed investment costs.

Given $S \in \mathcal{S}$, suppose that player j uses a mixed strategy s_j that plays the pure strategies $a_j^1, \ldots, a_j^t$ with probabilities $p_1, \ldots, p_t$. Then in $\Omega \cup S$, player j uses a mixed strategy that plays the pure strategies $\sigma_j \cup a_j^1, \ldots, \sigma_j \cup a_j^t$ with probabilities $p_1, \ldots, p_t$, where σ_j is the pure stragtegy used in Ω. We then obtain the following result, concerning any strategy set $S \in \mathcal{S}$.

Lemma 3.1. *Let γ be a submodular set function. Then for any $S \in \mathcal{S}$*

$$\bar{\gamma}(\Omega) \;\le\; \bar{\gamma}(S) + \sum_{i:\sigma_i \neq s_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) - \sum_{i:s_i \neq \sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$$

Proof. Observe that, by Lemma 2.2,

$$\begin{aligned} \bar{\gamma}(\Omega \cup S) \;&\le\; \bar{\gamma}(S) + \sum_{i:\sigma_i \neq s_i} \bar{\gamma}'_{\sigma_i}(S \cup \Omega^{i-1}) \\ &\le\; \bar{\gamma}(S) + \sum_{i:\sigma_i \neq s_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) \end{aligned}$$

In addition,

$$\bar{\gamma}(\Omega \cup S) \;=\; \bar{\gamma}(\Omega) + \sum_{i:s_i \neq \sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$$

Thus,

$$\bar{\gamma}(\Omega) \;\le\; \bar{\gamma}(S) + \sum_{i:\sigma_i \neq s_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) - \sum_{i:s_i \neq \sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$$

□

Now let us focus specifically on the case of Nash equilibria. We then obtain the following guarantee concerning the social value of a Nash equilibrium.

Theorem 3.2. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have*

$$\text{OPT} \leq 2\,\bar{\gamma}(S) - \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - \sum_{i:s_i\neq\sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$$

Proof. Observe that

$$\begin{aligned}
\sum_{i:\sigma_i\neq s_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) &\leq \sum_{i:\sigma_i\neq s_i} \max_{t_i\in\mathcal{S}_i} \bar{\gamma}'_{t_i}(S \oplus \emptyset_i) \\
&\leq \sum_{i:\sigma_i\neq s_i} \bar{\alpha}_i(S) && \text{[since } S \text{ is a Nash equilibrium]} \\
&\leq \bar{\gamma}(S) - \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) && \text{[by (1) and (2)]}
\end{aligned}$$

Note that Ω is a strategy set consisting of pure strategies. Therefore $\text{OPT} = \gamma(\Omega) = \bar{\gamma}(\Omega)$. So we have

$$\begin{aligned}
\text{OPT} &\leq \bar{\gamma}(S) + \sum_{i:\sigma_i\neq s_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i) - \sum_{i:s_i\neq\sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) \quad \text{[by Lemma 3.1]} \\
&\leq 2\bar{\gamma}(S) - \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - \sum_{i:s_i\neq\sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})
\end{aligned}$$

□

Observe that, for a general submodular function γ, the term $\sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)$ and/or the the term $\sum_{i:s_i\neq\sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1})$ may be negative. Thus, the social value of the Nash equilibrium is at least half the social optimal subject to a function-dependent additive term. As mentioned, this additive term often has a economic/social meaning. An alternative type of guarantee is also available. This result has clean implications in certain problems, for example, in the traffic routing problem of Section 7.

Theorem 3.3. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have*

$$2\bar{\gamma}(S) \geq \bar{\gamma}(\Omega \cup S) + \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i)$$

Proof.

$$\begin{aligned}
2\bar{\gamma}(S) &\geq \gamma(\Omega) + \sum_{i:s_i\neq\sigma_i} \bar{\gamma}'_{s_i}(\Omega \cup S^{i-1}) + \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \\
&= \bar{\gamma}(\Omega \cup S) + \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{s_i}(S \oplus \emptyset_i) \\
&= \bar{\gamma}(\Omega \cup S) + \sum_{i:s_i=\sigma_i} \bar{\gamma}'_{\sigma_i}(S \oplus \emptyset_i)
\end{aligned}$$

□

Theorem 3.4. *Let γ be a non-decreasing submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have*

$$\text{OPT} \leq 2\bar{\gamma}(S)$$

Proof. For non-decreasing, submodular functions the additive term in Theorem 3.2 is positive and, hence, we obtain a factor 2 guarantee. □

We remark that in some cases (depending upon the value of a measure of curvature of the non-decreasing submodular set function) Theorem 3.4 can be strengthened slightly; we omit the details. Note also that the social objective function in the marriage game is non-decreasing (the value of the game can not decrease with additional proposals) so we obtain a factor 2 guarantee for Nash equilibria in this game. Theorems 3.2 and 3.4 are both tight. We will give an example to show this in Section 6 when we discuss the competitive k-median problem.

4. Pure Strategy Nash Equilibria

Recall Theorem 2.1 which states that finite, non-cooperative, k-agent games have a Nash equilibrium. Unfortunately this is just an existence result and offers no help in actually finding Nash equilibria. In addition, the result just guarantees the existence of a mixed strategy Nash equilibrium. It is not the case that there need be pure strategy Nash equilibria; in fact, generally complex games (and many simple games) will not have a pure strategy Nash equilibria. The existence of pure strategy Nash equilibria is of interest for several reasons. In many practical situations, e.g. decisions concerning the location of facilities, agents are likely to adopt pure strategies. They are unlikely to chose one action amongst many on the basis of a coin toss. Furthermore, the strategy space of pure strategies is much smaller than the strategy space of mixed strategies. Thus, the discovery of pure strategy Nash equilibria may become a feasible. Moreover, given this smaller space, it is more reasonable to imagine that the agents can and will act in such a way as to generate a pure strategy Nash equilibria. In this section, we will show that any basic utility system has pure strategy Nash equilibria. We will also discuss how such equilibria may be realised in practice.

Theorem 4.1. *Take a valid utility system $(\gamma, \cup_i \alpha_i)$. If the utility system is basic then there are pure strategy Nash equilibria.*

Proof. Consider a directed graph D, each node of which corresponds to one of the possible pure strategy sets (i.e. action sets). There is an arc from node $\{a_1, a_2, \ldots, a_i, \ldots, a_k\}$ to node $\{a_1, a_2, \ldots, a'_i, \ldots, a_k\}$ if $\alpha_i(\{a_1, a_2, \ldots, a_i, \ldots, a_k\}) < \alpha_i(\{a_1, a_2, \ldots, a'_i, \ldots, a_k\})$, for some agent i. It follows that a node $\{a_1, a_2, \ldots, a_k\}$ in D corresponds to a pure strategy Nash equilibrium if and only if the node has out-degree zero. In particular, the system has a pure strategy Nash equilibrium if D is acyclic. We will show that for basic utility systems this is indeed the case.

Suppose D is not acyclic. Then take a directed cycle C in D. Suppose the cycle contains nodes corresponding to the action sets $A_0 = \{a_1^0, a_2^0, \ldots, a_k^0\}, A_1 = \{a_1^1, a_2^1, \ldots, a_k^1\}, \ldots, A_t = \{a_1^t, a_2^t, \ldots, a_k^t\}$ where $A_0 = A_t$. It follows that the action sets A_r and A_{r+1} differ in only the action of one agent, say agent i_r. Thus $a_i^r = a_i^{r+1}$ if $i \neq i_r$, and $\alpha_{i_r}(A_r) < \alpha_{i_r}(A_{r+1})$, that is

$$\alpha_{i_r}(\{a_1^r, a_2^r, \ldots, a_k^r\}) < \alpha_{i_r}(\{a_1^{r+1} = a_1^r, \ldots, a_{i_{r-1}}^{r+1} = a_{i_{r-1}}^r, a_{i_r}^{r+1}, a_{i_{r+1}}^{r+1} = a_{i_{r+1}}^r, \ldots, a_k^{r+1} = a_k^r\})$$

In particular, it must be the case that $\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) > 0$. We will obtain a contradiction by showing that, in fact, $\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) = 0$. Now $\alpha_{i_r}(A_{r+1}) = \gamma'_{a_{i_r}^{r+1}}(A_{r+1} \oplus \emptyset_{i_r})$ and $\alpha_{i_r}(A_r) = \gamma'_{a_{i_r}^r}(A_r \oplus \emptyset_{i_r})$. Thus

$$\begin{aligned}
\alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) &= \gamma'_{a_{i_r}^{r+1}}(A_{r+1} \oplus \emptyset_{i_r}) - \gamma'_{a_{i_r}^r}(A_r \oplus \emptyset_{i_r}) \\
&= (\gamma(A_{r+1}) - \gamma(A_{r+1} \oplus \emptyset_{i_r})) - (\gamma(A_r) - \gamma(A_r \oplus \emptyset_{i_r})) \\
&= (\gamma(A_{r+1}) - \gamma(A_r)) + (\gamma(A_r \oplus \emptyset_{i_r}) - \gamma(A_{r+1} \oplus \emptyset_{i_r})) \\
&= \gamma(A_{r+1}) - \gamma(A_r)
\end{aligned}$$

Here the last equality follows from the observation that $a_i^r = a_i^{r+1}$ if $i \neq i_r$. Then, since $A_0 = A_t$, we obtain

$$\begin{aligned}
\sum_{r=0}^{t-1} \alpha_{i_r}(A_{r+1}) - \alpha_{i_r}(A_r) &= \sum_{r=0}^{t} \gamma(A_{r+1}) - \gamma(A_r) \\
&= \gamma(A_t) - \gamma(A_0) \\
&= 0
\end{aligned}$$

□

Observe that Theorem 4.1 states not only that a pure strategy Nash equilibrium exists, but the proof also shows how one may be obtained. Specifically, if we start with any pure strategy set S (for example, $S = \{\emptyset_1, \emptyset_2, \ldots \emptyset_k\}$) and the agents sequentially alter their actions in order to maximise their own profits then we will automatically converge to a pure strategy Nash equilibrium. In addition, this is true even if the agents do not chose an optimal response at each step, but rather just chose any action that leads to an improvement in their private utility. So suppose that agents can quickly adapt their actions. Then pure strategy Nash equilibria can be generated just by the agents acting in any greedy fashion.

We note that for Theorem 4.1 we do require that the utility system be basic. For example, suppose we have a utility system $(\gamma, \cup_i \alpha_i)$ in which $\gamma(A) = M$, for some large constant M. Hence

g is a constant function and is, therefore, submodular. Consequently, we have $\gamma'_{a_i}(A \oplus \emptyset_i) = 0$. It follows that, if the system is not basic, the only constraints on the private utility functions are that $\sum_i \alpha_i(A) \leq M$, $\forall A \in \mathcal{A}$ and that $\alpha_i(A) \geq 0$, $\forall i$. However, this presents no real restriction on the game, other than that the private payoffs must be non-negative. It is, therefore, easy to give examples with no pure strategy Nash equilibria.

5. A Broader Framework

In this section we relax our third assumption, that is $\bar{\alpha}_i(S) \geq \bar{\gamma}'_{s_i}(S \oplus \emptyset_i)$. Instead we will consider the situation in which the private utility of an agent is comparable to the Vickrey utility with respect to that agent (loss in social utility that would result from the agent withdrawing from the game).

We say that $(\gamma, \cup_i \alpha_i)$ is a *(P,Q)-utility system* if, for some constants $P, Q > 0$,

$$\bar{\alpha}_i(S) \;\; \geq \;\; \frac{1}{P}\,\bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - Q \tag{5}$$

A (P,Q)-utility system is *(P,Q)-basic* if we have equality in condition (5): $\bar{\alpha}_i(S) = \frac{1}{P}\,\bar{\gamma}'_{s_i}(S \oplus \emptyset_i) - Q$. The system is *valid* if $\sum_i \bar{\alpha}_i(S) \leq \bar{\gamma}(S)$. Then we easily obtain the following results.

Theorem 5.1. *Let γ be a submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid (P,Q)-utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* OPT $\leq (1 + P)\,\bar{\gamma}(S) + (kQ - \sum_i \bar{\gamma}'_{s_i}(S \cup \Omega - s_i))$. □

Theorem 5.2. *Let γ be a non-decreasing, submodular set function. If $(\gamma, \cup_i \alpha_i)$ is a valid (P,Q)-utility system then for any Nash equilibrium $S \in \mathcal{S}$ we have* OPT $\leq (P + \delta(\gamma))\,\bar{\gamma}(S) + kQ$. □

Theorem 5.3. *Take a valid (P,Q)-utility system. If the system is (P,Q)-basic then there are pure strategy Nash equilibria.* □

6. The Competitive Facility Location and k-Median Problems

In this section we consider the facility location and k-median problems. First we will describe the problems and then introduce competitive versions of the problems. We will then show that these competitive problems fit into the framework given in the previous sections.

6.1. the base problems.

Both these facility location problems have the following form. We are given a bipartite graph $G = (W \cup U, E)$ with vertex partition W and U. The set W consists of locations at which facilities may be built. The set U consists of locations at which consumers are found. For clarity, we will refer to vertices in W as locations and the vertices in U as markets. In the base problems we have a single agent or monopolistic firm. The monopolist wishes to construct facilities at various locations in W in order to maximise its profits.

Each market u in U has an associated *value* π_u. A facility may be built at a location v for a *fixed cost* c_v. A facility at location v is able to service a market located at u for the *marginal cost* λ_{vu}. The *marginal profit* of the firm is its revenue minus its marginal costs. The *profit* of the firm is its marginal profit minus its fixed costs (i.e. revenue minus total costs). The *consumer surplus* is defined to be the total value minus total price. The *social surplus* is defined to be profits plus consumer surplus or, equivalently, total value minus total costs.

Let us examine these terms in more detail. Consider the revenue of the firm. This is just the sum of the prices it charges each market for servicing it. What will this price be, though, in the monopolistic case? Observe that consumers in market u have no choice but to be serviced by the monopolist. Their only constraint is that they will not pay more that π_u; thus, the firm will charge u a price $p_u = \pi_u$. It follows that consumer surplus is zero in the monopolist case. Thus a firm maximising profits is also, inadvertently, maximising the social surplus.

Observe that the firm will refuse to service a market u from a facility v if $\lambda_{vu} > \pi_u$. Thus a firm can always obtain a marginal revenue of zero with respect to each market. Thus our objective function will not be affected if we assume that our bipartite graph is complete and we have $\lambda_{vu} \leq \pi_u$ for each edge vu (that is setting $\lambda_{vu} = \pi_u$ where $\lambda_{vu} > \pi_u$ will not affect the outcome).

For the facility location problem, the firm may open whichever facilities it desires. So, formally, the *facility location problem* is

$$\max_{A \subseteq W} \mu(A) \quad = \quad \max_{A \subseteq W} \quad \left(\sum_u \max_{v \in A} (\pi_u - \lambda_{vu}) \quad - \quad \sum_{v \in A} c_v \right)$$

In the k-median problem the firm faces an additional constraint in that it can open at most k facilities. Formally, the *k-median problem* is

$$\max_{A \subseteq W, |A| \leq k} \mu(A) \quad = \quad \max_{A \subseteq W, |A| \leq k} \quad \left(\sum_u \max_{v \in A} (\pi_u - \lambda_{vu}) \quad - \quad \sum_{v \in A} c_v \right)$$

The performance of algorithms for these problems has been widely studied, (see, for example, [3],[11],[2],[5] and [1]). Note, it is often assumed that for the k-median problem there are no fixed costs i.e. $c_v = 0, \forall v$.

We also remark that, recently, the minimisation versions of both these problems have also received widespread attention (see, for example, [4] and [7]). The minimisation problems correspond to minimising the total costs of servicing all the markets. The broader economic viewpoint implied by the traditional maximisation problem, though, allows for very clean competitive formulations. It is these formulations that we will now introduce.

6.2. THE COMPETITIVE PROBLEMS.

The base problems correspond to the monopolistic situation. The corresponding competitive problem is as follows. Instead of a single monopoly, suppose we have k competing firms (or agents). In the *competitive facility location problem* the number of facilities each firm may open is unrestricted; whereas in the *competitive k-median problem* each firm may build at most one facility (in fact, our results hold for a more general problem in which firm i can open at most m_i facilities). We allow firms to build at the same location, but assume, however, that the costs differ for each firm. Thus firm i, $1 \leq i \leq k$, may build a facility at location v for a fixed cost c_v^i. In addition, the marginal cost of firm i servicing a market u from a facility at location v is λ_{vu}^i. Again, the value of market u is π_u.

The competitive situation differs markedly from the monopolistic case. Consider, for example, the pricing strategies of firms in non-competitive and competitive markets. We have seen that in the monopolistic case there is no consumer surplus; the monopoly gets all of the social surplus for itself. In a competitive market, though, firms have to compete for the market u. Let $\lambda_u^1, \lambda_u^2, \dots, \lambda_u^k$ be the lowest marginal costs with which the firms can supply market u, i.e.

$$\lambda_u^i = \min(\lambda_{v,u}^i : \text{firm } i \text{ has an open facility at } v)$$

and let $\lambda_u = \min_i \lambda_u^i$. What will happen in such a situation? Let $I_u^* = \{j : j = \text{argmin}_i \lambda_u^i\}$ be the collection of most competitive firms with respect to market u. Then, not surprisingly, a firm $i_u^* \in I_u^*$ will compete most efficiently and will thus service market u. However, the firm will not be able to charge π_u; instead, it will only be able to charge the marginal cost of the second most efficient firm. Thus u will pay a price of $p_u = \min_{i \neq i_u^*} \lambda_u^i$ in order to be serviced. If the firm i_u^* tries to charge more than this it will be under-cut by another firm. Since the price p_u may be less than π_u, positive consumer surpluses may now arise. Hence, the social surplus is indeed shared between the individual firms and the consumers; market u contributes $\pi_u - p_u$ to the consumer surplus and $p_u - \lambda_u$ to the marginal profits of the firm that services it. (It may be the case that multiple firms all have the lowest marginal costs with respect to a market u, that is $|I_u^*| \geq 2$. In such circumstances we will assume that customers in u randomly allocate their custom between these firms. The marginal profits for these firms will, though, be zero with respect to a market u, since they will compete away each others profits.)

Let $\Gamma_i = \{u : i \in I_u^*\}$ and $n_u = |I_u^*|$. Then, given a set of actions $A = A_1 \times A_2 \times \cdots \times A_k$ we have:

The profit of each firm i is

$$\begin{aligned}\omega_i(A) &= \sum_{u\in\Gamma_i}\frac{(p_u-\lambda_u^i)}{n_u}-\sum_{v\in a_i}c_v^i\\ &= \sum_{u\in\Gamma_i}(p_u-\lambda_u^i)-\sum_{v\in a_i}c_v^i\end{aligned}$$

The consumer surplus is

$$\begin{aligned}\zeta(A) &= \sum_u(\pi_u-p_u)\\ &= \sum_i\sum_{u\in\Gamma_i}\frac{(\pi_u-p_u)}{n_u}\end{aligned}$$

The social surplus is

$$\begin{aligned}\mu(A) &= \sum_u(\pi_u-\lambda_u^{i_u^*})-\sum_i\sum_{v\in a_i}c_v^i\\ &= \sum_i\sum_{u\in\Gamma_i}\frac{(\pi_u-\lambda_u^i)}{n_u}-\sum_i\sum_{v\in a_i}c_v^i\end{aligned}$$

So from a social viewpoint it would be best for a single authority to direct where each firm should locate in order to maximise the social surplus (utility). However, the firms themselves will choose strategies according to their own private profit (utility) functions. We next show, however, that these competitive formulations fit into the framework we have developed and, thus, we are able to obtain guarantees concerning the social performance on Nash equilibria in these facility location problems.

Before doing so we remark that it is common practice to present the facility location problem, as we have done, in terms of building facilities at specific locations in order to service markets. This, though, appears to be at odds with the statement that, from a practical point of view, our game-theoretic analysis is best suited to problems in which strategies are easy to change. Note, though, that we can view the problem in the following manner. Instead of facility location decisions we have fixed investment decisions. These fixed investments enable the firms to service various markets at specific marginal costs. Thus, from this wider perspective, the problem is that of making fixed investments in order to allow access to markets. From this perspective these facility location problems are very suitable for a game-theoretic analysis, as it is quite plausible that these investment decisions can be easily adapted.

6.3. THE SOCIAL PERFORMANCE OF NASH EQUILIBRIA IN THE FACILITY LOCATION PROBLEMS. First we need to show that we can formulate both the competitive facility location problem and the competitive k-median problem appropriately for our purposes.

Lemma 6.1. *The competitive facility location and k-median problems can be formulated in the action set framework.*

Proof. Consider first the competitive facility location problem. Recall that for our base problems we have a bipartite graph $G = (W \cup U, E)$. It follows that each agent i has a groundset $V_i = W$. Now, since a firm may open facilities at any set of locations, we have $\mathcal{A}_i = \{X : X \subseteq V_i\}$. Next consider the competitive k-median problem. Again, each agent i has a groundset $V_i = W$. Now, since each firm may open at most one facility we have $\mathcal{A}_i = \emptyset \cup \{v : v \in V_i\}$. $\square$

Next we need to show that our social utility (surplus) function is submodular.

Lemma 6.2. *The social surplus function μ is submodular.*

Proof. So

$$\mu(A) = \sum_u (\pi_u - \lambda_u^{i_u^*}) - \sum_i \sum_{v \in a_i} c_v^i = h(A) - g(A)$$

Now, clearly, $g(A) + g(B) = g(A \cap B) + g(A \cup B)$, for $A, B \subseteq V = \cup_i V_i$. So it suffices to show that h is submodular i.e. $\sum_u \lambda_u^{i_u^*}$ is supermodular. In what follows, we add an action set descriptor to distinguish between the four types of action set ($A, B, A \cap B$ and $A \cup B$). Let $i \in N_u(A \cup B)$. Without loss of generality, assume that $\lambda_u(A \cup B) = \lambda_{v_i,u}^i$ where $v_i \in A$. Then $\lambda_u(A \cup B) = \lambda_u(A)$. Clearly, however, $\lambda_u(A \cap B) \geq \lambda_u(B)$. It follows that $h(A) + h(B) \geq h(A \cap B) + h(A \cup B)$. $\square$

As mentioned, traditionally, the k-median problem is usually presented in the absence of fixed costs i.e. $c_v^i = 0$, $\forall i \, \forall v$. Such a formulation gives the following property.

Corollary 6.3. *In the absence of fixed costs, the social surplus function μ is non-decreasing.*

Proof. In the absence of fixed costs we have $g(A) = 0$, $\forall A \subseteq V$. Clearly h is a non-decreasing function, and hence μ is also non-decreasing. $\square$

Lemma 6.4. *The system $(\mu, \cup_i \omega_i)$ is a valid utility system. In particular, the utility system is basic.*

Proof. Recall that our private utility (profit) functions are

$$\omega_i(A) = \sum_{u \in \Gamma_i} (p_u - \lambda_u^i) - \sum_{v \in a_i} c_v^i$$

We now show that $(\mu, \cup_i \omega_i)$ is a basic utility system, that is

$$\mu'_{a_i}(A \oplus \emptyset_i) = \mu(A) - \mu(A \oplus \emptyset_i) = \sum_{u \in \Gamma_i} (p_u - \lambda_u^i) - \sum_{v \in a_i} c_v^i$$

The change in the social utility is the increase in the total marginal profits minus the increase in the total fixed cost, when agent i changes its action from the null action to action a_i. The increase in total marginal profits, though, is just the sum over all markets of the extra efficiency gained by i having action a_i. This, in turn, is the difference between the marginal costs of i, in those markets where it is the most efficient firm, and the marginal costs of the next most efficient firm. This is just $\sum_{u \in \Gamma_i}(p_u - \lambda_u^i)$. Clearly the total change in the fixed costs is $\sum_{v \in a_i} c_v^i$, as required. Hence, the system $(\mu, \cup_i \omega_i)$ is a basic utility system. By Theorem 2.5, the utility system is also valid. □

We are now in the position to apply Theorems 3.2 and 3.4. If we denote by $FC(S)$ and $MP(S)$ the expected fixed costs and expected marginal profits, respectively, associated with a solution $S \in \mathcal{S}$, then

Theorem 6.5. *For the competitive k-median and facility locations problems, any Nash equilibrium $S \in \mathcal{S}$ satisfies*

$$\text{OPT} \;\leq\; 2\,\bar{\mu}(S) + FC(S) \;=\; \bar{\mu}(S) + MP(S) + \bar{\zeta}(S)$$

Proof. From Theorem 3.2 we have $\bar{\mu}(\Omega) \leq 2\bar{\mu}(S) - \sum_{i:s_i=\sigma_i} \bar{\mu}'_{s_i}(S \oplus \emptyset_i) - \sum_{i:s_i \neq \sigma_i} \bar{\mu}'_{s_i}(\Omega \cup S^{i-1})$. On the addition of an extra firm to the game, the social surplus can deteriorate by at most the fixed costs incurred by the new firm. Thus the additive term is upper bounded by $FC(S)$. The result then follows from the observation that $MP(S) - FC(S) + \bar{\zeta}(S) = \bar{\mu}(S)$. □

We will see that Theorem 6.5 is tight. Let us first comment briefly upon its implications. The theorem tell us that our guarantee is good when either the fixed costs or the marginal profits plus consumer surplus induced by the solution S are small compared to OPT. Conversely, if the fixed costs and marginal profits plus consumer surplus are both large then the overall social performance may be very poor. Such a situation may arise in industries in which there are high start-up costs combined with markets that contain a collection of highly valuable customers. As a result, firms may over-supply the valuable customers (at the expense of less valuable customers) leading to a wasteful duplication of services. Such examples are common in the high-tech industry where the occurrence of high initial costs often allows a firm access to lucrative markets. A less obvious example is the health industry. Here there are very high fixed and initial costs in the actual provision of health care, and also in the associated revenue collection system (e.g. insurance companies, HMOs, finance departments in hospitals, etc). In addition, the market also contains many highly valuable customers from both the private sector (companies with large workforces) and the public sector (government supported Medicare patients). In contrast, there is a large class of less valuable customers. The resulting social inefficiencies are illustrated by the large number of uninsured citizens, as well as the duplication of services.

In the absence of fixed costs, we obtain, from Theorem 3.4

Theorem 6.6. *Consider the competitive k-median problem in the absence of fixed costs. Then any Nash equilibrium $S \in \mathcal{S}$ satisfies*

$$\text{OPT} \;\le\; (1+\delta(\mu))\,\bar{\mu}(S) \;\le\; 2\,\bar{\mu}(S) \qquad \Box$$

6.4. PURE STRATEGY NASH EQUILIBRIA AND FACILITY LOCATION.

Observe that both facility location problems have the desirable property that they possess pure strategy Nash equilibria. This follows from Theorem 4.1 and Lemma 6.4.

Theorem 6.7. *For the competitive facility location and k-median problems there exist pure strategy Nash equilibria.* $\Box$

6.5. TIGHT EXAMPLES.

We previously claimed that Theorems 3.2 and 3.4 were both tight. We will now prove this using the following example concerning the competitive 2-median problem shown in Figure 1. Let there be two agents with two possible locations, v_1 and v_2, at which to locate; we use the superscripts 1 and 2 to distinguish between the respective copies of v_i, $i \in \{1,2\}$. In addition, there are four markets u_1, u_2, u_3 and u_4. The value of each market is one. All marginal costs are 1, except for the six (represented by labelled edges) shown in the figure. In addition, there are no fixed costs. Thus the social surplus function is non-decreasing and submodular.

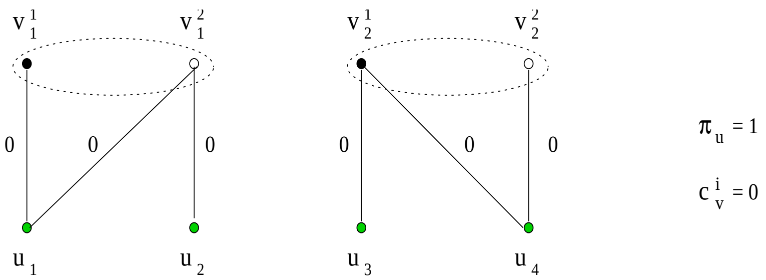

FIGURE 1. A tight example.

The optimal strategy set is $\Omega = \{v_2^1, v_1^2\}$, i.e. firm 1 should use the pure strategy (action) of locating at v_2, whilst firm 2 should use the pure strategy of locating at v_1. Such a strategy pairing will give a social surplus of 4.

We remark that the strategy set Ω is also a Nash equilibrium. However, there are other Nash equilibria. Consider though the pure strategy set $S = \{v_1^1, v_2^2\}$. It is easy to verify that this is also Nash equilibrium. Each firm has a private profit of 1 under S, and if they change their strategy (whilst the other sticks with its strategy) they still receive a profit of 1. The social surplus of this strategy set is 2. Thus, the social value of this Nash equilibrium is a factor 2 off that of the optimal solution.

Next we show that Theorem 3.2 is also tight. To do this we, again, use an example from the 2-median problem, see Figure 2. There are still two possible locations, v_1 and v_2. However, we have the following fixed costs. Firm 1 can locate at v_1^1 for a cost 1, but can locate at v_2^1 for nothing; firm 2 can locate at v_2^2 for a cost 1, but can locate at v_1^2 for nothing. These fixed costs are shown boxed in Figure 2. We have two markets u_1 and u_2, both of which have a value 1. All marginal cost are also 1 except for the four shown, which are zero.

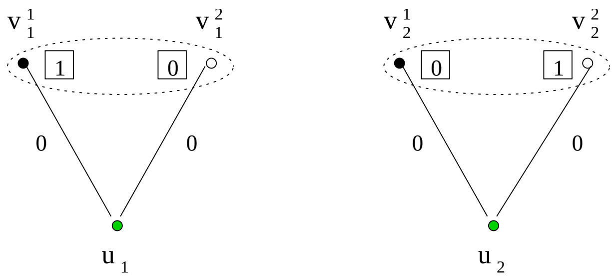

FIGURE 2. A tight example.

The optimal strategy set is $\Omega = \{v_2^1, v_1^2\}$, i.e. firm 1 should use the pure strategy (action) of locating at v_2, whilst firm 2 should use the pure strategy of locating at v_1. Such a strategy pairing will give a social surplus of 2 (and private profits of one each).

Again, the strategy set Ω is also a Nash equilibrium. However is is easy to check that the pure strategy set $S = \{v_1^1, v_2^2\}$ is also Nash equilibrium. This Nash equilibrium has a social surplus of 0. The fixed costs of the solution, though, are 2. Similarly, the marginal profits of the solution are also 2. So Theorem 6.5 and, hence, Theorem 3.2 are tight.

7. The Selfish Traffic Routing Problem

In this section we consider the problem of routing traffic in a network. Congestion in the network causes delays and is costly for individual agents and society as a whole. It would help, therefore, if the traffic could be directed by a single authority. However, it is individual agents who make their own routing decisions. Thus the problem appears suitable for analysis via our techniques. In particular, here we sketch how a maximisation version of the selfish routing problem of Roughgarden and Tardos [16] fits into our framework. They considered the following network routing problem. There is a directed network $G = (V, A)$ and k source-destination vertex pairs, $\{s_1, t_1\}, \ldots, \{s_k, t_k\}$ (note that we do not require k to be large). The collection of paths from s_i to t_i is denoted by $\mathcal{P}_i$ with $\mathcal{P} = \cup_i \mathcal{P}_i$. A flow is a function $f : \mathcal{P} \to \mathbb{R}^+$; for a fixed flow f, we have $f_a = \sum_{P \in \mathcal{P}: a \in P} f_P$. Now $f = \cup_i f_i$ where f_i is a flow from s_i to t_i. We will abuse our notation slightly and also denote by f_i the value of the flow f_i; given the context this should not cause any confusion.

Each arc $a \in A$ has a load-dependent *latency function*, denoted by $l_a(f)$. The latency of a path P with respect to a flow f is defined as the sum of the latencies of the edges in the path, denoted by $l_P(f) = \sum_{a \in A} l_a(f_a)$. The latency with respect to an agent i is $l_i(f) = \sum_{P_i \in \mathcal{P}_i} l_{P_i}(f) f_{P_i}$. The latency $l(f)$ of a flow f is the total latency incurred by f i.e.

$$l(f) \;=\; \sum_{a \in A} l_a(f_a)\, f_a \;=\; \sum_{P \in \mathcal{P}} l_P(f)\, f_P \;=\; \sum_i l_i(f)$$

In [16] the social objective is to minimise the total latency, given that a flow of value r_i must be routed from s_i to t_i. The private objective of an agent i is to minimise its own latency i.e. $l_i(f)$.

We consider a maximisation version of this problem. Each agent may route a flow of weight at most r_i from s_i to t_i. Associated with each source-destination $\{s_i, t_i\}$ pairing is a value π_i that signifies the revenue (utility) from routing one unit of flow from s_i to t_i. However, we still associate with a routing the latency-based cost. Thus, a flow f that successfully routes f_i units of flow from s_i to t_i will induce a profit to agent i of $\zeta_i(f) = \pi_i\, f_i - l_i(f)$. Hence, the social objective is to maximise the function

$$\kappa(f) \;=\; \sum_i \zeta_i(f) \;=\; \sum_i \pi_i\, f_i - l_i(f)$$

and agent i seeks to maximise the private objective function ζ_i. We will now show that this problem also fits into our framework. To do this we will discretise the problem by assuming that flow may be sent only in whole unit increments; for this problem it is not difficult to generalise the results to continuous space.

Lemma 7.1. *The routing problem can be formulated in the action set framework.*

Proof. The action space $\mathcal{A}_i$ of agent i consists of any flow f_i of value at most r_i from s_i to t_i. We now show how this fits into our framework. For each agent i we have a collection of paths $\mathcal{P}_i$ from s_i to t_i. The agent assigns a weight to each path $p_i \in \mathcal{P}_i$. Let the groundset V_i consist of r_i copies of each path p_i i.e. $p_i^1, \ldots, p_i^{r_i}$. Here the choice of p_r^t correspond to the routing of t units of flow on path p_i.

We may allow an agent to select multiple copies of a path. In such a circumstance only the action corresponding to the copy with the greatest amount of flow is implemented. (Alternatively, we may restrict the action space of agent i to allow for the choice of at most one copy of each path p_i). Note that if no copy of p_i is chosen then no flow is sent along that path. $\square$

Now consider that latency functions $l_a(f)$. We will assume that these functions are non-negative, non-decreasing and convex. Note that these assumptions correspond to some natural properties of traffic systems. The non-decreasing property implies that the costs incurred increase as the volume of the traffic increases; the convexity property implies that the additional costs incurred (by adding an additional unit of traffic) increase as the volume of the traffic increases. Observe that convexity implies that the latency functions are supermodular when restricted to our discretised space. It follows easily that

Lemma 7.2. *For the selfish routing problem, the social objective function κ is submodular.* $\square$

Lemma 7.3. *For the selfish routing problem, the system (κ, ζ_i) is a valid utility system.*

Proof. We will show, for each agent i, that $\zeta_i(f) \geq \kappa'_{f_i}(f \oplus \emptyset_i)$. Now

$$\begin{aligned}
\kappa'_{f_i}(f \oplus \emptyset_i) &= \kappa(f) - \kappa(f - f_i) \\
&= \sum_j (f_j\,\pi_j - l_j(f)) - \sum_{j:j\neq i} (f_j\,\pi_j - l_j(f - f_i)) \\
&= f_i\,\pi_i - l_i(f) + \sum_{j:j\neq i} (l_j(f - f_i) - l_j(f)) \\
&\leq f_i\,\pi_i - l_i(f) \\
&= \zeta_i(f)
\end{aligned}$$

Thus (κ, ζ_i) is a utility system. We have already seen that $\kappa(f) = \sum_i \zeta_i(f)$ and, thus, the utility system is valid. $\square$

So we then obtain the following guarantees.

Theorem 7.4. *For the selfish routing problem, any Nash equilibrium $S \in \mathcal{S}$ satisfies*

$$\text{OPT} \leq 2\,\bar{\kappa}(S) - \sum_{i:s_i=\sigma_i} \bar{\kappa}'_{s_i}(S \oplus \emptyset_i) - \sum_{i:s_i\neq\sigma_i} \bar{\kappa}'_{s_i}(\Omega \cup S^{i-1}) \qquad \square$$

Thus we obtain a factor 2 guarantee if, for example, $\bar{\kappa}'_{s_i}(S \oplus \emptyset_i), \bar{\kappa}'_{s_i}(\Omega \cup S^{i-1}) \geq 0,\ \forall i$. An alternative guarantee follows from Theorem 3.3. This compares the value of a Nash equilibrium S against the social value of a particular solution, $S + \Omega$, that routes twice as much traffic.

Theorem 7.5. *For any Nash equilibrium $S \in \mathcal{S}$, we have*

$$2\bar{\kappa}(S) \geq \bar{\kappa}(\Omega \cup S) + \sum_{i:s_i=\sigma_i} \bar{\kappa}'_{\sigma_i}(S \oplus \emptyset_i) \geq \bar{\kappa}(S + \Omega) \quad \square$$

A result of this flavour also follows from the work of [16]; the social value of a Nash equilibrium is at least the social value of the optimal solution that routes twice as much traffic when the all the rewards π_i are halved.

If κ is non-decreasing (hence, it is always in the interest of agent i to route all r_i units of flow), then from Theorem 3.4 we obtain

Theorem 7.6. *If κ is non-decreasing then, for the selfish routing problem, any Nash equilibrium $S \in \mathcal{S}$ satisfies*

$$\text{OPT} \leq 2\,\bar{\kappa}(S) \qquad \square$$

8. Polynomial Time Considerations

Our discussion regarding pure strategy Nash equilibria touched upon the importance of speed considerations in the strategy determination. We discuss this in more detail in this section. Let us measure the size of the problem input in terms of the size of the groundsets V_i, $1 \leq i \leq k$. It would be useful if we obtained a Nash equilibria in polynomial time in the problem size. Two factors are important here:

(i) Bounding the number of times an agent changes strategy before a Nash equilibria is obtained.

(ii) Bounding the time an agent takes to decide upon a strategy.

How to bound the number of iterations required before convergence to a Nash equilibria is an important open question. In the presence of pure Nash equilibria, as we have seen, the overall size of the state space gives one upper bound. We note, however, that good guarantees may be obtained within a constant number of iterations (we only need each agent to change strategies a constant number of times). That is, solutions that arise long before we reach a Nash equilibria also provide good guarantees. Thus, although these solutions may not be stable, they do give good performance. We omit the details here.

Regarding the second factor, if the size of the action space $\mathcal{A}_i$ of agent i is polynomial in $|V_i|$, then the agent can easily find its best strategy in polynomial time. However, the action space $\mathcal{A}_i$ may be as large as $2^{|V_i|}$. Thus in some circumstances it may not be possible to find an optimal strategy quickly. It may, though, be possible to obtain approximately optimal strategies in polynomial time. We will show that the use of approximation algorithms by the agents in their strategy determination does lead to guarantees on the social performance of Nash equilibria. We have one difficulty to overcome though. The use of approximately optimal strategies is not consistent with the concept of a Nash equilibria. That is, approximately optimal strategies are *not* the optimal best response strategies required by Nash equilibria. Thus, we are really using *approximate Nash equilibria.* They are equilibria in the sense that no agent can find (by whatever methods they are using) a better alternative strategy in polynomial time.

So suppose that each agent has access to an approximation algorithm at each stage. Let these algorithms have an approximation guarantee of ξ, say. Then, Theorem 3.2, Theorem 3.3 and Theorem 3.4 apply (with slightly weaker guarantees) to approximate Nash equilibria. For example, if our social utility function is non-decreasing, we have the following theorem.

Theorem 8.1. *Let γ be a non-decreasing, submodular set function, and $(\gamma, \cup_i \alpha_i)$ be a valid utility system. If the agents can generate ξ-approximate solutions, then for any approximate Nash* *equilibrium* $S \in \mathcal{S}$ *we have*

$$\text{OPT} \;\leq\; (\xi + \delta(\gamma))\,\bar{\gamma}(S) \;\leq\; (\xi + 1)\,\bar{\gamma}(S) \qquad \Box$$

For an example consider the case of matroids. A *matroid* $\mathcal{T}$ is a family of subsets of V such that

(i) $\emptyset \in \mathcal{T}$.

(ii) If $Y \in \mathcal{T}$ and $X \subseteq Y$, then $X \in \mathcal{T}$.

(iii) If $X, Y \in \mathcal{T}$ and $|X| < |Y|$, then $\exists y \in Y - X$ such that $X \cup \{y\} \in \mathcal{T}$.

Fisher, Nemhauser and Wolsey [6] gave a simple 2-approximation algorithm for the problem of maximising a non-decreasing, submodular function over a matroid. Thus, if each action set $\mathcal{A}_i$ is a matroid then we have

Corollary 8.2. *Let* γ *be a non-decreasing, submodular set function, and* $(\gamma, \cup_i \alpha_i)$ *be a valid utility system. If each* $\mathcal{A}_i$ *is a matroid, then we obtain an approximate Nash equilibrium* $S \in \mathcal{S}$ *with*

$$\text{OPT} \;\leq\; (2 + \delta(\gamma))\,\bar{\gamma}(S) \;\leq\; 3\,\bar{\gamma}(S) \qquad \Box$$

9. Multiple-Item Auctions

Consider the following class of auction: there is one seller (auctioneer) with a set J of n different items, and a set of k potential buyers (agents) who have a private valuation for each subset of items. One form of auction within this class is *combinatorial auctions*. These are auctions in which agents may make bids on subsets of items (*combinatorial bids*), rather than just bids on individual items. There is a very large literature on combinatorial auctions; see de Vries and Vohra [17] for a survey. The following are factors which the seller may wish to consider when designing an auction structure in which to sell the items.

(1) Simplicity: the rules of the auction should be easily understood.

(2) Fairness: agents need to believe that the rules of the auction are fair.

(3) Speed: the auction should not take too long to complete.

(4) Efficiency: the seller may wish to allocate the items to maximise the social value.

(5) Revenue: the seller wants to maximise the total revenue it receives from the auction.

Note that goals 4) and 5) may not be compatible. Hence, in this section we will focus on goals 1) to 4). We will also be concerned with the case in which the private valuation function v_i, for each buyer i, is submodular i.e. the marginal valuations are non-increasing. Recently, Lehmann, Lehmann and Nisan [9] considered the *allocation problem* induced by this framework. There, a single authority wishes to find an allocation of optimal efficiency (social value). They present a polynomial time algorithm that produces an allocation with social value at least one half that of the optimal solution, provided that the agents valuations are submodular. Their approach is as follows. The authority knows (or has access to) each agents valuation function. The authority then greedily assigns one item at a time, say in the order $1, 2, \ldots, n$. Let $V(j)$ the be value of the allocation after the jth item is assign. Item $j+1$ is then assigned to the agent so as to maximise $V(j+1) - V(j)$. That is, item $j+1$ is assigned to the agent with the highest *marginal valuation* for the item, given the allocation of items $1, 2, \ldots, j$. It can be shown that the allocation produced by such a process is, indeed, at least half optimal.

Again, our interest is in the competitive situation in which the seller and buyers all seek to maximise their own utility. We present a simple class of *multi-round auction* that is guaranteed to produce an allocation within a factor 2 of optimal, despite the valuation functions being private knowledge and with the sellers and buyers acting in a selfish manner. Moreover, the allocation procedure of [9] can easily be implemented within this class of auction.

9.1. THE RULES OF THE AUCTION.

We now give the rules of the auction. In the first round, the seller sets a price p_j for each item j in the auction. Each buyer then states which items it is willing to purchase at these prices. If more than one agent accepts the price p_j then in the next round the auctioneer will raise the price (by any amount it chooses) of item j. If no agent accepts the price p_i then in the next round the auctioneer will lower the price (by any amount it chooses) of item j. After each round the auctioneer announces *provisional winners* for each item. The provisional winner of an item will be randomly selected from amongst those agents that have the highest bid for the item. The announcement of provisional winners tells the agents who will win the items if the auction were to terminate at that time. This information allows the agents to make bids with the knowledge of whether their bids from previous rounds have been "accepted". *Provisionally winning* bids are considered binding and cannot be withdrawn. A provisionally winning bid for an item only ceases to be of interest after a higher bid for that item has been made. However, in future rounds, agents may ignore any bids they made that were not provisionally winning. The auction terminates when there is exactly one bidder for each every item, and no agent wishes to change its bid (that is, bid for a set of items that is a superset of its current set of winning items).

[We remark that it is important that provisionally winning bids cannot be withdrawn; if bids can be withdrawn then the results (that will follow) regarding polynomial time convergence are lost. It should be noted, however, if this auction did allow the withdrawal of bids then we would actually converge to an optimal allocation. To see this, suppose that $\{T_1, T_2, \ldots, T_k\}$ is an optimal allocation but $\{S_1, S_2, \ldots, S_k\}$ is the solution produced by the auction, with termination item prices $\{p_1, p_2, \ldots, p_n\}$. Now for each agent, $v_i(S_i) - \sum_{j \in S_i} p_j \geq v_i(T_i) - \sum_{j \in T_i} p_j$ otherwise agent i would have changed its bid to T_i. Summing over all agents we have

$$\begin{aligned}
\sum_i v_i(S_i) - \sum_i \sum_{j \in S_i} p_j &\geq \sum_i v_i(T_i) - \sum_i \sum_{j \in T_i} p_j \\
\sum_i v_i(S_i) - \sum_{j \in J} p_j &\geq \sum_i v_i(T_i) - \sum_{j \in J} p_j \\
\sum_i v_i(S_i) &\geq \sum_i v_i(T_i)
\end{aligned}$$

Thus, $\{S_1, S_2, \ldots, S_k\}$ is an optimal allocation.]

9.2. PERFORMANCE GUARANTEES.

It is clear that this auction does satisfies the goal of simplicity. It also satisfies the goal of fairness since the highest bidder for an item wins it (with possibly a random choice in the case of a tie). In contrast, note that in combinatorially auctions it is not always clear to the agents that items are allocated in a "fair" manner. Next, we consider the issue of efficiency. In order to do this we need to examine the actions of agents in such an auction. Faced with a set of prices how do the agents react. To begin with we will assume that the agents act in a *myopically rational* manner, see [13], that is, they make a best response to the current prices and allocation. Hence an agent "bids" on all the items in a subset that maximise its utility given the stated prices (this includes all it bids that are currently provisionally winning bids). Later we will show that our performance guarantees still hold even when the agents are allowed to make *locally myopically optimal* bids (to be defined). This generalisation is useful as it is easy for the agents to find locally myopically optimal bids, whereas obtaining the myopically optimal bid may take exponential time.

Note that the valuation functions of an agent are submodular. Thus, since we have a fixed price per item, the private utility functions (i.e. private valuation minus auction price) are also submodular. Now, at a given stage in the auction, suppose that agent i has provisionally winning bids for a set S_i of the items at the current prices. Then, since bids cannot be withdrawn, in the next stage the agents must optimise with respect to the groundset $J - S_i$. That is, the agent must look to bid on other items given that it has already bid for S_i. For example, when considering a set $X \subseteq J - S_i$, the agent must evaluate the set by considering $v_i(X \cup S_i)$ not $v_i(X)$, since the agent is already contracted to buy S_i (the agent stops being contacted to buy an item j in S_i only if the price p_j rises in a later phase and another agent accepts the new price but agent i does not).

Let us consider the utility to agent i if it is allocated the set S_i in the auction. The agent pays a price $p(S_i) = \sum_{j \in S_i} p_j$ for the set of items and, thus, receives a private utility of $u_i(S_i) = v_i(S_i) - p(S_i)$. So the goal of the each agent is to maximise its private utility. The social utility denoted by $\eta(S)$ is just the sum of the values of the sets in the allocation produced by the bidding strategies $S = \{S_1, \ldots, S_k\}$. Observe that $\eta(S)$ also equals the sum of the private utilities plus the revenue from the auction (that is, the utility of the auctioneer).

Lemma 9.1. *Take a Nash equilibrium S, then for any agent i, we have*

$$u_i(S) = \eta(S) - \eta(S \oplus \emptyset_i) - p(S_i)$$

Proof. Note that a Nash equilibrium corresponds to a completed auction. Thus, each agent i is sold the subset S_i of items it bid for. Recall that we may view only the provisional winning bids as being binding and hence we may assume that no other agent has a binding bid on any of these items. Therefore $\eta(S) - \eta(S \oplus \emptyset_i) = v_i(S_i)$. To see this, note that if agent i were able to withdraw its bids then the social value of the auction would fall by $v_i(S_i)$ since no other bidders has a binding for those items at the current prices. Now $u_i(S) = v_i(S_i) - p(S_i)$ and the lemma follows. □

In order to apply our results we must ensure that our auction can be implemented in our framework. This, though, is easy. Assume that the auctioneer has set the prices, then an action of agent i is just a subset of the groundset i.e. which item prices the agent accepts. For the purposes of analysis we may "pretend" that there are multiple copies of each item, and that each agent receives a copy of an item if the agent accepts an item price. This allows us to assign a social value to outcomes like $\eta(\Omega \cup S)$ which were used in the previous proofs (the social value is just the sum of the values of the set of items assigned to each agent). Note that for any real auction solutions though, for example S and Ω, there must be exactly one winning bid for each item.

So the auction does fit into our framework, but it is not immediately obvious that we can now apply Theorem 3.2. This is because we have $u_i(S) = \bar{\eta}'_{s_i}(S) - p(S_i)$ rather than $u_i(S) = \bar{\eta}'_{s_i}(S)$. Fortunately, however, we also have $\eta(S) = \sum_i u_i(S) + \sum_i p(S_i)$ rather than $\eta(S) = \sum_i u_i(S)$. It is easy to check that these differences cancel each other out in the proof, and so it follows that Theorem 3.2 does indeed hold in this auction problem.

We will make the standard assumption that there are zero disposable costs. Thus, the private valuation functions v_i are non-decreasing. Since we evaluate (possibly non-feasible) solutions from the "multiple-copies" viewpoint that all bids are accepted, it follows that η is a non-decreasing function. Hence by Theorem 3.4 we obtain

Theorem 9.2. *The social value of any auction solution S satisfies*

$$\text{OPT} \;\leq\; (1+\delta(\eta))\,\bar{\eta}(S) \;\leq\; 2\,\bar{\eta}(S) \qquad \square$$

It is also easy to check that Theorem 5.3 still applies given Lemma 9.1. Thus we have

Theorem 9.3. *The auction has pure strategy Nash equilibria.* □

So given that the agents bid in a myopically rational manner, the auction produces an allocation with efficiency within a factor 2 of the most efficient allocation.

9.3. FAST IMPLEMENTATIONS.

As mentioned, in the case of auctions fast implementation is very important. In practice, this means take the time required for the auction should be polynomial in the number of items. Here we will outline some of the issues involved. Firstly, how long does each round take? Until now we have assumed that each agent bids in a myopically rational manner and, thus, maximises a submodular objective function at each round. This may be too time consuming for our purposes. However, the performance guarantees hold even when the agents bid in a simple locally optimal (greedy) manner at each stage. Agent i when faced with a set of prices $\{p_1, p_2, \ldots, p_n\}$ greedily chooses a subset S_i as follows. Initially $S_i = \emptyset$. Add to S_i an item j such that $u_i(S_i \cup \{j\}) > u_i(S_i)$, that is $v_i(S_i \cup \{j\}) > v_i(S_i) + p_j$, then repeat. If no such item exists then stop.

We now discuss why such a strategy is locally myopically optimal. First note that such a bidding strategy restricts agent i to bid for a set of items S_i with the property that $u_i(T_i) \geq 0$, for all $T_i \subseteq S_i$. We call this the *risk-free* property; we say that a bidding strategy that is not risk-free is *risky*. To see why the agents (without any information regarding the private valuations of the other bidders) will wish to adopt a *risk-free* bidding strategy, suppose instead that an agent adopts a risky strategy. It is then easy to provide the other agents with private valuations that ensure that the agent receives a set X that induces a negative utility. Thus, without any information regarding the private valuations of the other bidders, the agents will wish to restrict their attention to risk-free bidding strategies. In addition, the bidding strategy given above, also ensures that an agent bids for a maximal risk-free sets. These maximal sets are locally myopically optimal; to see this, suppose we have a risk-free set S_i that is not maximal, then there is an item j such that $u_i(T_i \cup \{j\}) > u_i(T_i)$, for all $T_i \subseteq S_i$.

It can easily be shown that if the agents make maximal risk-free bids then a factor two approximation guarantee is also obtained. We sketch a proof. Such bids ensures that each agent has positive utility; thus, $\eta(S) \geq \text{REV}$, where REV is the revenue the auctioneer receives from the auction. Moreover, it can also easily be shown that $\eta(S) \geq \text{OPT} - \text{REV}$. To see this, suppose that $\{T_1, T_2, \ldots, T_k\}$ is an optimal allocation but $\{S_1, S_2, \ldots, S_k\}$ is the solution produced by the auction, with termination item prices $\{p_1, p_2, \ldots, p_n\}$. Now for each agent, $v_i(S_i) - \sum_{j \in S_i} p_j \geq v_i(S_i \cup T_i) - \sum_{j \in S_i \cup T_i} p_j$ otherwise there is an item $j \in T_i - S_i$ such that $u_i(S_i \cup \{j\}) > u_i(S_i)$ and agent i would not have bid for S_i. Summing over all agents we have

$$\begin{aligned}
\sum_i v_i(S_i) - \sum_i \sum_{j\in S_i} p_j &\geq \sum_i v_i(S_i \cup T_i) - \sum_i \sum_{j \in S_i \cup T_i} p_j \\
\sum_i v_i(S_i) &\geq \sum_i v_i(S_i \cup T_i) - \sum_i \sum_{j\in T_i} p_j \\
\sum_i v_i(S_i) &\geq \sum_i v_i(T_i) - \sum_i \sum_{j\in T_i} p_j \\
\eta(S) &\geq \text{OPT} - \text{REV}
\end{aligned}$$

The result then follows.

We can polynomially bound, using standard bisection method techniques, the number of rounds required to complete the auction easily. For example, we now show how the allocation procedure of [9] can be implemented by such an auction. The auctioneer initially announces a set of prices $\{p_1, p_2, \ldots, p_n\} = \{V, V, \ldots, V\}$ (where V is an upper bound on the value any bidder attaches to any single item) and then changes the prices of each item, in turn, until there is exactly one bidder for the item. Note that, when an item is considered the items that still have price V will have no bidders. It follows that the agent that has the greatest marginal valuation for that item (given the current allocation induced by the items that have already been considered) will be the agent that makes the highest bid on the item. Note that, by submodularity, no agent will want to bid for an item in a later round after it has been considered (even though such bids are allowed). Thus we obtain assignment procedure of [9], and the implementation time is polynomial in the number of items. To see this observe that, by bisection techniques, the number of rounds required to complete the auction is at most $O(n \log V)$.

Remark An obvious question here is whether better performance guarantees can be obtained in auctions which allow combinatorial bidding.

Acknowledgments. I am especially grateful to Tim Roughgarden and Nicolas Stier for comments and suggestions that greatly improved this paper. I am also very grateful to Vahab Mirrokni, Andreas Schulz, Bruce Shepherd and Santosh Vempala for interesting discussions on this work. Finally, I would like to thank Eva Tardos for suggesting auctions as a possible application of the method.

References

[1] A. Ageev and M. Sviridenko, "An 0.828-approximation algorithm for the uncapacitated location problems", *Discrete Applied Mathematics*, **93**, pp289-96, 1999.

[2] M. Conforti and G. Cornuéjols, "Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the Rado-Edmonds theorem", *Discrete Applied Mathematics*, **7**, pp251-274, 1984.

[3] G. Cornuéjols, M. Fischer and G. Nemhauser, "Location of bank accounts to optimize float: an analytic study of exact and approximation algorithms", *Management Science*, **23**, pp789-810, 1977.

[4] M. Charikar, S. Guha, E. Tardos and D. Shmoys, "A constant-factor approximation algorithm for the k-median problem", *Proceedings of the* 31^{st} *Annual Symposium on the Theory of Computing*, pp1-10, 1999.

[5] G. Cornuéjols, G. Nemhauser and L. Wolsey, "The uncapacitated facility location problem", in P. Mirchandani and R. Francis eds., *Discrete Location Theory*, Wiley, New York, pp119-171, 1990.

[6] M. Fisher, G. Nemhauser and L. Wolsey, "An analysis of approximations for maximising submodular set functions II", *Math. Programming Study*, **8**, pp73-87, 1978.

[7] K. Jain and V. Vazirani, "Primal-dual approximation algorithms for metric facility location and k-median problems", *Proceedings* 40^{th} *Annual Symposium on the Foundations of Computer Science*, pp2-13, 1999.

[8] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria", *Proceedings of the* 16^{th} *Annual Symposium on Theoretical Aspects of Computer Science*, pp404-413, 1999.

[9] B. Lehmann, D. Lehmann and N. Nisan, "Combinatorial auctions with decreasing marginal utilities", *ACM Conference on Electronic Commerce*, 2001.

[10] J. Nash, "Non-cooperative games", *Annals of Mathematics*, **54**, pp268-95, 1951.

[11] G. Nemhauser and L. Wolsey, "Best algorithms for approximating the maximum of a submodular set function", *Mathematics of Operations Research*, **3(3)**, pp177-188, 1978.

[12] C. Papadimitriou, "Algorithms, games, and the internet", *Proceedings of the* 33^{rd} *Annual Symposium on the Theory of Computing*, 2001.

[13] D. Parkes and L. Ungar, "Iterative combinatorial auctions: theory and practice', *Proceedings of the* 17^{th} *National Conference on Artificial Intelligence*, pp74-81, 2000.

[14] M. Rothkopf, A. Pekec, R. Harstad, "Computationally manageable combinational auctions", *Management Science*, **44**, pp1131-47, 1998.

[15] T. Roughgarden, "Designing networks for selfish users is hard", *Proceedings of the* 42^{st} *Annual Symposium on the Foundations of Computer Science*, pp473-81, 2001.

[16] T. Roughgarden and E. Tardos, "How bad is selfish routing?", *Journal of the ACM*, **49(2)**, pp236–259, 2002.

[17] S. de Vries and R. Vohra, "Combinatorial auctions: a survey", *INFORMS Journal on Computing*, **15**, 2003.