

Distributed Welfare Games

Jason R. Marden

Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309, jason.marden@colorado.edu

Adam Wierman

Department of Computer Science, California Institute of Technology, Pasadena, CA 91125, adamw@caltech.edu

The assignment of utility functions is a central component of a game theoretic approach to distributed resource allocation. The goal is to assign each decision making entity a utility function such that the resulting game possesses a host of desirable properties including scalability, tractability, locality, and existence and efficiency of pure Nash equilibria. In this paper we formally study this question of utility design. We introduce a class of games termed distributed welfare games as our platform for utility design. By viewing utility design in a similar fashion to cost sharing, we identify several utility design methodologies that guarantee desirable game properties irrespective of the specific application domain. Lastly, we illustrate the results in this paper on two well studied resource allocation problems: the graph coloring problem and the vehicle target assignment problem.

Key words: resource allocation; game theory; distributed control

History: A few of the results in this paper appeared in (Marden and Wierman (2008)), which did not include full proofs.

1. Introduction

Resource allocation problems are of fundamental importance in many applications. Traditionally, researchers have focused on developing centralized algorithms to determine efficient allocations (Feige and Vondrak (2006), Ageev and Sviridenko (2004), Ahuja et al. (2003)). However, in many modern applications, these centralized algorithms are neither applicable or desirable. One concrete example is the problem of sensor coverage where the goal is to allocate a fixed number of sensors across a given mission space so as to maximize the probability of detecting a particular event (Li and Cassandra (2005)). In this case, a centralized algorithm requires that a central authority maintains complete knowledge of the environment and can communicate directly with each sensor during the entire mission. Both requirements may be unrealistic in large and/or hostile environments. The same issues arise in many computer network resource allocation problems, e.g., wireless access point assignment (Kaumann et al. (2007)), congestion control (Garcia et al. (2000), Akella et al. (2002)), and wireless power management (E. Campos-Nañez (2008), Li and Cassandra (2005)). There are also many examples outside of computer systems, e.g., minimizing aggregate congestion in a transportation system. In this setting, a global planner does not have the authority to assign drivers to roads; rather, a global planner must entice drivers appropriately, possibly through taxation, to settle on a desirable allocation (Sandholm (2002)).

Recently, there has been surge of research aimed at understanding the possibility of decentralizing (localizing) decisions in resource allocation problems (Kaumann et al. (2007), Mhatre et al. (2007), Komali and MacKenzie (2007), Zou and Chakrabarty (2004), Srivastava et al. (2005), E. Campos-Nañez (2008)). One approach for accomplishing this is to model the resource allocation problem as a noncooperative game where the players, e.g, the sensors in the sensor coverage problem or the drivers in the transportation systems, independently pursue their own objectives which may or may not be in conflict with other players. There are many advantages to this game-theoretic architecture including robustness to failures and environmental disturbances, reducing communication requirements, improving scalability, etc (Arslan et al. (2007)).

Several challenges arise when seeking to design and implement such a distributed system. One of the primary challenges is the assignment of utility functions where the goal is to assign utility functions such that the resulting game has desirable properties (Arslan et al. (2007), Marden et al. (2007a)). For example,

in the sensor coverage problem a design goal could be to assign each sensor a utility function that depends only on *local* information and guarantees that the resulting Nash equilibria are efficient with respect to the global objective.

Noncooperative resource allocation has recently been applied to a number of computer network applications, e.g., Akella et al. (2002), Kaumann et al. (2007), Mhatre et al. (2007), Komali and MacKenzie (2007), Zou and Chakrabarty (2004), E. Campos-Nañez (2008). However, the design of utility functions is highly dependent on the application domain primarily because of the lack of utility design methodologies in the existing literature. For example, E. Campos-Nañez (2008) focuses on efficiently managing a tradeoff between energy usage and sensing capability in sensor networks while Komali and MacKenzie (2007) focuses on topology control in ad-hoc networks. The respective utility designs are as independent as the problem domains. However, the analysis is virtually identical as both papers investigate existence and efficiency of (pure) Nash equilibria among other things. Are application specific designs necessary or is it possible to develop utility design methodologies for any domain that would provide these guarantees? Establishing such methodologies is of paramount importance for the propagation of game theory as a valuable tool for distributed resource allocation.

The goal of this paper is to establish a general framework for (i) investigating the feasibility of noncooperative resource allocation and (ii) designing desirable utility functions. To that end, we consider a class of resource allocation problems with separable objective/welfare functions. Such a class of problems induces a natural notion of locality that we would like to satisfy in our utility design. By local, we mean that each player's utility should depend only on the selected resources and the other players that selected the same resources. To formally study utility design, we introduce a class of game that we refer to as *distributed welfare games* (see Section 3.2). In a distributed welfare game, each player's utility is defined as some fraction of the welfare garnered at each of the player's selected resources. Therefore, designing utility functions is equivalent to defining a distribution rule that depicts how the welfare garnered from a specific allocation is distributed to the players. The primary goal is to design distribution rules for distributed welfare games that guarantee the existence and efficiency of pure Nash equilibrium which represent individually agreeable allocations. Other properties that we investigate include locality, informational dependencies, budget balance, and whether the resulting game is a potential game.

While the motivation for locality and information dependencies is clear, the motivation for budget balance and potential games requires further explanation. Budget balance imposes a constraint on the sum of the individual players utility functions. In particular, the sum of the players utility functions needs to equal the global welfare. Budget balanced utility functions are important in many application domains where there is a cost or revenue that needs to be completely absorbed by the participating players, e.g., network formation (Chen et al. (2008)) or content distribution (Goemans et al. (2004)). In more traditional engineering systems where such a constraint is not relevant it turns out that having budget balanced utility functions is still desirable. This stems from the recent results in (Roughgarden (2009), Gairing (2009)) which show that having budget balanced utility functions can help provide robust efficiency guarantees. As for potential games, if the distribution rule results in the formulation of a potential game, then one can appeal to a variety of distributed algorithms to ensure that the players converge to a Nash equilibrium (Young (1998), Monderer and Shapley (1996a), Marden et al. (2009)).

The first set of results (Section 4) illustrate that cost sharing methodologies can be used effectively for utility design. In particular, we identify two cost sharing methodologies, the *marginal contribution* and the *Shapley value*, that provide valuable methodologies for utility design. The marginal contribution distribution rule systematically provides local utility functions that always guarantee the existence of a Nash equilibrium irrespective of the application domain. The Shapley value distribution rule systematically provides local utility functions that are budget balanced and always guarantee the existence of a Nash equilibrium irrespective of the application domain. However, computing the Shapley value is often intractable. These results highlight an inherent tension between budget balance utility functions and tractability.

Our second set of results seek to overcome the limitations of the Shapley value by establishing tractable budget balanced distribution rules (Section 5). We identify three sufficient conditions on distribution rules,

see Conditions 5.1–5.3, that guarantee the existence of a Nash equilibrium in any distributed welfare game where players are restricted to selecting a single resource. These sufficient conditions can be viewed from two perspectives. The first perspective is as a check for whether a given set of distribution rules guarantees the existence of an equilibrium. The second perspective, which is our motivation for this work, is as a design guideline for distribution rules, i.e., if a global planner can design a distribution rule to satisfy these conditions then an equilibrium is guaranteed to exist. We demonstrate that using these conditions as a guideline for the design of distribution rules leads to the construction of desirable distribution rules in the vehicle target assignment problem, where by desirable we mean that the distribution rule is budget balanced, guarantees the existence of an equilibrium in all considered games, and is tractable.

Our third set of results pertain to the efficiency of the resulting Nash equilibria. We measure the efficiency of a Nash equilibrium using the well known measures: *price of anarchy* and *price of stability*. The price of anarchy (stability) is defined as the worst-case ratio between the global welfare evaluated at the worst (best) Nash equilibrium and the optimal welfare. These measures have been studied extensively in several application domains (Vetta (2002), Johari and Tsitsiklis (2004), Nissan et al. (2007)). In general, the price anarchy in distributed welfare games can be arbitrarily close to 0; however, when restricting attention to submodular welfare functions, which is a common attribute in many resource allocation problems (Vetta (2002), Krause and Guestrin (2007)), we can develop distribution rules that obtain a welfare within 1/2 of that of the optimal assignment. Furthermore, we tighten this price of anarchy bound in a variety of settings. This compares favorably with the best known results of centralized approximations for resource allocation problems with submodular welfare functions, which guarantee welfare within $1 - 1/e \approx 0.6321$ of the optimal (Feige and Vondrak (2006), Ageev and Sviridenko (2004), Ahuja et al. (2003)). Surprisingly, this comparison demonstrates that the inefficiency resulting from localizing decisions in resource allocation problems is relatively small when the welfare functions are submodular.

Lastly, we end the paper in Section 7.2 with an illustration of how the theory developed in this paper applies to two well studied resource allocation problems: the graph coloring problem and the vehicle target assignment problem. This section highlights that these methodologies can be applied to several alternative problem domains to construct desirable utility functions.

It should be noted that this paper predominantly focuses on equilibrium behavior in distributed welfare games. An alternative question that is of equal importance is understanding how players reach an equilibrium in a distributed fashion. While not focusing on this question in detail, we illustrate the applicability of the theory of *learning in games* (Young (1998), Monderer and Shapley (1996a), Marden et al. (2009, 2007b,a)) as a distributed control mechanism for coordinating group behavior. For example, if a distributed welfare game constitutes a potential game, then a global planner can appeal to a variety of distributed learning algorithms of varying complexity to guarantee that the group behavior converges to a Nash equilibrium (Young (1998), Monderer and Shapley (1996a), Marden et al. (2009, 2007b,a,c)). For a survey of how the learning in games literature and the question of utility design we study in this paper can be combined to develop distributed control algorithms and protocols see Gopalakrishnan et al. (2010).

2. Background: Finite Strategic-Form Games

We consider a finite strategic-form game with n -player set $N := \{1, \dots, n\}$ where each player $i \in N$ has a finite action set \mathcal{A}_i and a utility function $U_i : \mathcal{A} \rightarrow \mathbb{R}$ where $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ denotes the set of action profiles. For an action profile $a = (a_1, a_2, \dots, a_n) \in \mathcal{A}$, let a_{-i} denote the profile of player actions *other than* player i , i.e.,

$$a_{-i} = \{a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n\}.$$

With this notation, we will sometimes write a profile a of actions as (a_i, a_{-i}) . Similarly, we may write $U_i(a)$ as $U_i(a_i, a_{-i})$.

We focus on analyzing equilibrium behavior in such games. A well-known equilibrium concept that emerges in noncooperative games is that of a *pure Nash equilibrium*. An action profile $a^* \in \mathcal{A}$ is called a *pure Nash equilibrium* if for all players $i \in N$,

$$U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*). \quad (1)$$

A pure Nash equilibrium represents a scenario for which no player has a unilateral incentive to deviate. We will henceforth refer to a pure Nash equilibrium as simply an equilibrium.

We consider the class of potential games (Monderer and Shapley (1996b)) which, in the context of this paper, imposes a restriction on the class of admissible utility functions. In a potential game, the change in a player's utility that results from a unilateral change in strategy equals the change in a global potential function. Specifically, there is a potential function $\phi : \mathcal{A} \rightarrow \mathbb{R}$ such that for every player $i \in N$, for every $a_{-i} \in \mathcal{A}_{-i}$, and for every $a'_i, a''_i \in \mathcal{A}_i$,

$$U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i}) = \phi(a'_i, a_{-i}) - \phi(a''_i, a_{-i}).$$

When this condition is satisfied, the game is called a potential game with the potential function ϕ . It is easy to see that in potential games any action profile maximizing the potential function is a pure Nash equilibrium, hence every potential game possesses at least one such equilibrium.

3. Our Model

Game theory is emerging as a popular tool for distributed control of multi-agent systems. Utilizing game theory for distributed control requires modeling the interactions of the multi-agent system as a noncooperative game. The central component of this model is assigning a utility function U_i to each agent. The goal is to design utility functions such that the resulting game possess an array of desirable properties including existence and efficiency of equilibria, tractability, locality, budget balance, among others. To formally study the relationship between utility design and the resultant game properties we introduce a class of resource allocation problems with separable welfare functions as our platform for utility design. This class of resource allocation problem imposes a natural notion of locality which we seek to satisfy in our utility design. Understanding utility design for this class of resource allocation problems is of fundamental importance for understanding utility design in more general resource allocation problems.

3.1. Resource Allocation Problems

We consider the class of resource allocation problems where there exists a set of players N and a finite set of resources \mathcal{R} that are to be shared by the players. Each player $i \in N$ is assigned an action set $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$ where $2^{\mathcal{R}}$ denotes the power sets of \mathcal{R} ; therefore, a player may have the option of selecting multiple resources. A key feature of resource allocation problems is the existence of a global welfare function $W : \mathcal{A} \rightarrow \mathbb{R}$ that the system designer seeks to optimize. In this paper, we restrict our attention to the class of *separable* welfare functions of the form

$$W(a) = \sum_{r \in \mathcal{R}} W_r(\{a\}_r)$$

where $W_r : 2^N \rightarrow \mathbb{R}^+$ is the welfare function for resource r and $\{a\}_r := \{i \in N : r \in a_i\}$ is the set of players using resource r in the allocation a . Note that when restricting attention to separable welfare functions, the welfare generated at a particular resource depends only on which players are currently using that resource. While separable welfare functions cannot model the global objective for all resource allocation problems, they can model the global objective for several important classes of resource allocation problems including routing over a network (Roughgarden (2005)), vehicle target assignment problem (Murphey (1999a)), sensor coverage (Marden and Wierman (2008), Li. and Cassandras (2005)), content distribution (Goemans

et al. (2004)), graph coloring (Panagopoulou and Spirakis (2008)), network coding (Marden and Effros (2009)), among many others.

Throughout this paper we frequently restrict our attention to *submodular* welfare functions. A welfare function W_r is submodular if for any player sets $S \subseteq T \subseteq N$

$$W_r(S \cup i) - W_r(S) \geq W_r(T \cup i) - W_r(T)$$

for any $i \in N$. Submodularity corresponds to a notion of decreasing marginal returns and is a common feature of many objective functions for engineering applications ranging from content distribution (Goemans et al. (2004)) to coverage problems (Krause and Guestrin (2007)).

3.2. Distributed Welfare Games

We introduce the class of *distributed welfare games* to formally study the role of utility design in resource allocation problems with separable welfare functions. In a distributed welfare game, each agent's utility is defined as some fraction of the welfare garnered at each resource the agent is using. More formally, the utility of an agent for an action $a_i \subseteq \mathcal{R}$ is defined as

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} f_r(i, \{a\}_r) \quad (2)$$

where $f_r : N \times 2^N \rightarrow R$ is referred to as the *distribution rule* at resource r . We define a distributed welfare game by the tuple $G = \{N, \mathcal{R}, \{\mathcal{A}_i\}_{i \in N}, \{W_r\}_{r \in \mathcal{R}}, \{f_r\}_{r \in \mathcal{R}}\}$. For brevity, we will often omit the subscripts on the sets and denote a game as purely $G = \{N, \mathcal{R}, \{\mathcal{A}_i\}, \{W_r\}, \{f_r\}\}$

A few comments are in order regarding the structure imposed on the utility functions in (2). First, this structure imposes a natural notion of locality as each agent's utility function only depends on the resources the agent selected and the other agents that selected the same resources. Secondly, defining a distribution rule for each resource $\{f_r\}$ results in a well defined game irrespective of the structure of the players action sets $\{\mathcal{A}_i\}$. This allows for a degree of *scalability* in utility design as we can remove the dependence between utility design, or equivalently distribution rule design, and the structure of the players' action sets.

4. Utility Design for Distributed Welfare Games

In this section we explore several approaches for utility design in distributed welfare games. We discuss these design methodologies in the form of distribution rules as opposed to utility functions for a more direct presentation. Many of these approaches are derived from methodologies in the cost sharing literature, e.g., Shapley value, as this problem can be viewed in a similar context to a cost sharing problem (Young (1994), Shapley (1953), Hart and Mas-Colell (1989)). While cost sharing methodologies can be effective as distribution rules there are several issues that limit their applicability.

4.1. Performance Criteria

In this section we formally define our performance criteria before proceeding with the discussion on distribution rules. We assume throughout that the set of players N , resources \mathcal{R} , and welfare functions $\{W_r\}$ are all known; however the action sets $\{\mathcal{A}_i\}$ are unknown. We require that the distribution rule f_r for each resource $r \in \mathcal{R}$ can *only* depend on the welfare function W_r . Therefore distribution rules are completely defined using only local information. We gauge the quality of a design by the following metrics. As we will see, designing distribution rules to satisfy the following metrics is non-trivial and in many cases it is impossible.

Existence and efficiency of equilibrium: Does the distribution rule $\{f_r\}$ guarantee the existence and efficiency of an equilibrium irrespective of the structure of the players action sets $\{\mathcal{A}_i\}$?

Potential game: Does the distribution rule $\{f_r\}$ guarantee that the resulting game is a potential game irrespective of the structure of the players action sets $\{\mathcal{A}_i\}$?

Budget balance: A distribution rule $\{f_r\}$ is budget balanced if for any resource $r \in \mathcal{R}$ and any player set $S \subseteq N$, $\sum_{i \in S} f_r(i, S) = W_r(S)$.

Informational dependencies: This metric seeks to measure the informational dependencies between a distribution rule f_r and the associated the welfare function W_r . For any resource $r \in \mathcal{R}$, player $i \in N$ and player set $S \subseteq N$ such that $i \in S$, we expand the notation of a distribution rule from $f_r(i, S)$ to $f_r(i, S; *)$ to explicitly highlight the information dependencies, denoted by $(*)$, needed to compute the distribution to player i . We categorize informational dependencies as defined below:

- *High:* The distribution to player i is conditioned on the welfare for all subsets of S , i.e., $f_r(i, S; \{W_r(T)\}_{T \subseteq S})$. This means that the distribution rule requires complete knowledge of the structural form of the welfare function. For large player sets, this class of distribution rules may be intractable.

- *Medium:* The distribution to player i could incorporate information about marginal contributions to player set S , i.e., $f_r(i, S; \{W_r(T)\}_{T \in Z})$ where $Z = \{S, \{S \setminus j\}_{j \in S}\}$.

- *Low:* The distribution to player i could incorporate information about the current welfare to the player set S , i.e., $f_r(i, S; W_r(S))$.

4.2. Equally Shared Utilities

Suppose the welfare from each resource is divided equally amongst the players that selected the resource, i.e., for any set of players $S \subseteq N$

$$f_r^{\text{ES}}(S) = \frac{W_r(S)}{|S|}. \quad (3)$$

We refer to (3) as the equally shared utility design. It is straightforward to verify that this rule is budget balanced and has a low information dependency. However, in general such a design cannot guarantee the existence of an equilibrium as the following example illustrates. For alternative examples, see Arslan et al. (2007).

EXAMPLE 1 (EQUALLY SHARED UTILITIES). Consider a two player distributed welfare game with player set $N = \{1, 2\}$, resources $\mathcal{R} = \{r_1, r_2\}$, actions set $\mathcal{A}_1 = \mathcal{A}_2 = \mathcal{R}$, utility functions of the form (3), and a separable welfare function as illustrated below. Note that the game does not possess an equilibrium.

		Player 2		
		\emptyset	r_1	r_2
Player 1	\emptyset	0	4	1
	r_1	6	6	7
	r_2	5	9	10
		Welfare		

		Player 2		
		\emptyset	r_1	r_2
Player 1	\emptyset	0, 0	0, 4	0, 1
	r_1	6, 0	3, 3	6, 1
	r_2	5, 0	5, 4	5, 5
		Payoffs		

One problem with the equally shared utility design is that players' utility functions are not *aligned* with their contribution to the global welfare. However, if players are *anonymous* with regards to their impact on the global welfare function, i.e., for any resource $r \in \mathcal{R}$ and any player sets $S, T \subseteq N$ such that $|S| = |T|$ the welfare satisfies $W_r(S) = W_r(T)$, then the equally shared utilities in (3) guarantee the existence of an equilibrium. We refer to such welfare functions as anonymous.

PROPOSITION 1. Consider any distributed welfare game G with anonymous resource specific welfare functions. If the distribution rule satisfies (3), then an equilibrium is guaranteed to exist.

Proof: With a slight abuse of notation, let $W_r : \{0, 1, \dots\} \rightarrow R_+$ represent the anonymous welfare functions. Define $|a|_r := |\{a\}_r|$ as the number of players that chose resource r in the allocation a . It is straightforward to show that any distributed welfare game with anonymous players is a congestion game (Rosenthal (1973), Monderer and Shapley (1996b)), with the following specification:

- (a) resources: \mathcal{R} ,
- (b) cost functions: $c_r(k) = \frac{W_r(k)}{k}$, $k > 0$, where k is the number of players utilizing resource r , and
- (c) utility functions: $U_i(a) = \sum_{r \in a_i} c_r(|a|_r)$.

Any congestion game is a potential game with potential function

$$\phi(a) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{|a|_r} c_r(k).$$

Therefore, an equilibrium is guaranteed to exist in such a game.

□

4.3. Marginal Contribution Utilities

Consider an rule which distributes welfare according to a player's marginal contribution to the global welfare. In this setting, the distribution rule takes on the form

$$f_r^{\text{MC}}(i, S) = W_r(S) - W_r(S \setminus \{i\}) \quad (4)$$

for each player set $S \subseteq N$. This design is sometimes referred to as the *wonderful life utility* (WLU) (Wolpert and Tumor (1999)). It is well known that distributing the welfare as in (4) results in a potential game with potential function W ; hence any action profile that maximizes the global welfare is an equilibrium. However, other equilibria may also exist under the wonderful life utility design.

There are two limitations of the the marginal contribution utility design. First, each player needs to be able to compute his marginal contribution to the welfare in order to evaluate his utility, i.e., it has a medium informational dependency. Second, the wonderful life utility may distribute more (or less) welfare than is gathered; hence, it is not budget balanced. It remains an open question as to whether the second limitation can be addressed utilizing a similar informational requirement.

4.4. The Shapley Value

While the WLU guarantees the existence of an equilibrium in all settings, it may distribute more or less reward than the welfare garnered. It turns out that we can rectify this problem by distributing welfare according to a player's Shapley value (Shapley (1953), Hart and Mas-Colell (1989), Haeringer (2006)). That is, for any resource $r \in \mathcal{R}$ and player set $S \subseteq N$ the distribution rule takes on the form

$$f_r^{\text{SV}}(i, S) = \sum_{T \subseteq S \setminus \{i\}} \frac{|T|!(|S| - |T| - 1)!}{|S|!} (W_r(S \cup \{i\}) - W_r(T)). \quad (5)$$

Utilizing the Shapley value as in (5) requires a high informational dependency; however, it rectifies the budget balance problems associated with the WLU as shown in the following proposition.

PROPOSITION 2. *Consider any distributed welfare game G . If the distribution rule takes on the form in (5), then the ensuing game is budget balanced. Furthermore, it is a potential game with the following potential function $\phi^{\text{SV}} : \mathcal{A} \rightarrow \mathbb{R}$*

$$\phi^{\text{SV}}(a) := \sum_{r \in \mathcal{R}} \sum_{S \subseteq \{a\}_r} \frac{1}{|S|} \left(\sum_{T \subseteq S} (-1)^{|S|-|T|} W_r(T) \right). \quad (6)$$

Proof: We prove this proposition using the potential function derived in Hart and Mas-Colell (1989). First, we express the Shapley value distribution rule in (5) as a weighted sum of unanimity games (Hart and Mas-Colell (1989), Haeringer (2006)) which takes on the form

$$f_r^{\text{SV}}(i, S) = \sum_{T \subseteq S: i \in T} \frac{1}{|T|} \left(\sum_{R \subseteq T} (-1)^{|T|-|R|} W_r(R) \right). \quad (7)$$

Let $\alpha_r(T) := \sum_{R \subseteq T} (-1)^{|T|-|R|} W_r(R)$ and $\phi_r^{\text{SV}} : 2^N \rightarrow \mathbb{R}$ be the resource specific potential function (Hart and Mas-Colell (1989))

$$\phi_r^{\text{SV}}(a) := \sum_{T \subseteq \{a\}_r} \frac{\alpha_r(T)}{|T|}.$$

Let $a \in \mathcal{A}$ be any allocation and $a_i^0 = \emptyset$ be the null action for player i . Using (5), the marginal utility of player i is

$$\begin{aligned} U_i(a) - U_i(a_i^0, a_{-i}) &= \sum_{r \in a_i} f_r^{\text{SV}}(i, \{a\}_r), \\ &= \sum_{r \in a_i} \left(\sum_{T \subseteq \{a\}_r: i \in T} \frac{\alpha_r(T)}{|T|} \right), \\ &= \sum_{r \in a_i} \left(\sum_{T \subseteq \{a\}_r} \frac{\alpha_r(T)}{|T|} - \sum_{T \subseteq \{a\}_r \setminus \{i\}} \frac{\alpha_r(T)}{|T|} \right), \\ &= \sum_{r \in a_i} (\phi_r^{\text{SV}}(a) - \phi_r^{\text{SV}}(a_i^0, a_{-i})), \\ &= \phi^{\text{SV}}(a) - \phi^{\text{SV}}(a_i^0, a_{-i}). \end{aligned}$$

Therefore, for any player i , actions $a'_i, a''_i \in \mathcal{A}_i$, and allocation $a_{-i} \in \mathcal{A}_{-i}$

$$U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i}) = \phi^{\text{SV}}(a'_i, a_{-i}) - \phi^{\text{SV}}(a''_i, a_{-i}).$$

□

It is also worth noting that this potential function could be computed recursively. For any resources $r \in \mathcal{R}$ define $\phi_r^{\text{SV}}(\emptyset) := 0$. One can recursively evaluate ϕ_r^{SV} according to the following equation: for any $S \subseteq N$,

$$\phi_r^{\text{SV}}(S) = \frac{1}{|S|} \left[W_r(S) + \sum_{i \in S} \phi_r^{\text{SV}}(S \setminus \{i\}) \right].$$

There are two limitations of the Shapley value utility design that may prevent it from being applicable. First, there is a high informational requirement as each player must be able to compute his marginal contribution to all action profiles in order to evaluate his utility. Second, in general computing a Shapley value is intractable in games with a large number of players. This is highlighted explicitly in either (5) or (8) where computation of the Shapley value requires a weighted summation over all subsets of players. However, it should be noted that this computational cost is lessened dramatically if there are a limited number of distinct ‘‘classes’’ of players, see Conitzer and Sandholm (2004). For example, if players are anonymous then the Shapley value is equivalent to the equal share distribution rule in (3).

4.5. The Weighted Shapley Value

A generalization of the Shapley Value is the *weighted Shapley value* (Shapley (1953), Hart and Mas-Colell (1989), Haeringer (2006)). Define $w_i \in \mathbb{R}_+$ as the weight of player i . Let $w := \{w_i\}_{i \in N}$ be the associated weight vector. For the weighted Shapley value, the distribution rule takes on the following form: for any resource $r \in \mathcal{R}$ and player set $S \subseteq N$

$$f_r^{\text{WSV}}(i, S) := \sum_{T \subseteq S: i \in T} \frac{w_i}{\sum_{j \in T} w_j} \left(\sum_{R \subseteq T} (-1)^{|T|-|R|} W_r(R) \right). \quad (8)$$

Note that the Shapley value distribution rule in (5) is recaptured if $w_i = 1$ for all players $i \in N$. We state the following proposition without proof to avoid redundancy.

PROPOSITION 3. *Consider any distributed welfare game G . If the distribution rule takes on the form in (8), then the ensuing game is budget balanced. Furthermore, it is a (weighted) potential game.*

The weighted Shapley value does not result in as clean a closed form expression for the potential function as the Shapley value in (6). However, as with the Shapley value, the potential function can be computed recursively and is of the form

$$\phi^{\text{WSV}}(a) := \sum_{r \in \mathcal{R}} \phi_r^{\text{WSV}}(a),$$

where $\phi_r^{\text{WSV}}(a)$ is the resource specific potential function, $\phi_r^{\text{WSV}}(\emptyset) := 0$, and for any subset $S \subseteq N$ (Hart and Mas-Colell (1989))

$$\phi_r^{\text{WSV}}(S) = \frac{1}{\sum_{i \in S} w_i} \left[W_r(S) + \sum_{i \in S} w_i \phi_r^{\text{WSV}}(S \setminus \{i\}) \right].$$

Finally, note that the weighted Shapley utility design suffers from the same drawbacks as the Shapley value utility design.

4.6. Comparison of distribution rules

To this point we have surveyed four distribution rules motivated by methodologies from the cost sharing literature. We end this section by summarizing the advantages and disadvantages of these rules. Table 1 illustrates the tradeoff between desirable features of a distribution rule and the computational and informational requirement needed to obtain such a rule. The only setting where a distribution rule achieves all of the desired properties is the case where players are anonymous. In this setting, the Shapley value is equivalent to equally share distribution rule in (3). In the next section, our goal will be to develop other distribution rules that attain (nearly) all of our desired properties.

Distribution Rule	Existence of Equilibrium	Potential Game	Budget Balanced	Tractable	Informational Requirement
Equally Shared (anonymous)	yes	yes	yes	yes	Low
Equally Shared	no	no	yes	yes	Low
WLU	yes	yes	no	yes	Medium
Shapley	yes	yes	yes	no	High
Weighted Shapley	yes	yes	yes	no	High

Table 1 Summary of Distribution Rules for Distributed Welfare Games

5. Single Selection Distributed Welfare Games

The results of the previous section suggest that designing budget balanced distribution rules that always guarantee the existence of an equilibrium requires a high information requirement and is often intractable. In this section we explore whether this apparent tradeoff is a limitation of cost sharing methodologies or utility design in general. To study this question we focus on a simplified setting where players are only allowed to select a single resource, $\mathcal{A}_i = \mathcal{R}$, as opposed to multiple resources, $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. In this restricted setting, we identify three sufficient conditions for distribution rules that guarantee the existence of an equilibrium. These sufficient conditions can be viewed from two perspectives. The first perspective is as a check for whether a given set of distribution rules guarantees the existence of an equilibrium. The second perspective, which is our motivation for this work, is as a design guideline for distribution rules, i.e., if a global planner can design a distribution rule to satisfy these conditions then an equilibrium is guaranteed to exist. In Section 7.2 we demonstrate that using these conditions as a guideline for the design of distribution rules leads to the construction of desirable distribution rules in the well studied vehicle-target assignment problem. By desirable, we mean that the distribution rule is budget balanced, guarantees the existence of an equilibrium in all considered games, and has a low informational requirement.

5.1. Sufficient Conditions for Existence of an Equilibrium

In this section we identify three sufficient conditions on distribution rules that guarantee the existence of an equilibrium in any single selection resource allocation game. These sufficient conditions translate to pairwise comparisons of players' distributed shares.

CONDITION 5.1. Let $i, j \in N$ be any two players. If

$$f_r(i, S \cup \{i, j\}) > f_r(j, S \cup \{i, j\})$$

for some resource $r \in \mathcal{R}$ and player set $S \subseteq N \setminus \{i, j\}$, then

$$f_{\bar{r}}(i, \bar{S} \cup \{i, j\}) \geq f_{\bar{r}}(j, \bar{S} \cup \{i, j\})$$

for any resource $\bar{r} \in \mathcal{R}$ and player set $\bar{S} \subseteq N \setminus \{i, j\}$. For this situation we say that player i is stronger than player j . Furthermore, note that strengths are transitive, i.e., if player i is stronger than player j who is stronger than player k , then player i is also stronger than player k .

CONDITION 5.2. If player i is stronger than player j , then for any resource $r \in \mathcal{R}$ and player set $S \subseteq N \setminus \{i, j\}$ we have

$$f_r(i, S \cup \{i\}) \geq f_r(i, S \cup \{i, j\}).$$

CONDITION 5.3. If player i is stronger than player j , then for any resource $r \in \mathcal{R}$ and player set $S \subseteq N \setminus \{i, j\}$ we have

$$\frac{f_r(j, S \cup \{j\})}{f_r(i, S \cup \{i\})} \geq \max_{\bar{r} \in \mathcal{R}} \frac{f_{\bar{r}}(j, S \cup \{i, j\})}{f_{\bar{r}}(i, S \cup \{i, j\})}.$$

THEOREM 1. Consider any distributed welfare game G . If for all players $i \in N$ the action set is $\mathcal{A}_i = \mathcal{R}$ and the distribution rule $\{f_r\}$ satisfies Conditions 5.1, 5.2, and 5.3 then an equilibrium exists.

Proof: We begin by renumbering the players in order of strengths with player 1 being the strongest player. This is possible because of Condition 5.1.

We construct an equilibrium by letting each player select a resource one at a time in order of strength. The general idea of the proof is that once a player selects a resource, the player will never seek to deviate regardless of the other players selections. First, player 1 selects the resource $r^{(1)}$ according to

$$r^{(1)} \in \arg \max_{r \in \mathcal{R}} f_r(1, \{1\}) \tag{9}$$

Denote the action profile $a^{(1)} = (r^{(1)}, \emptyset, \dots, \emptyset)$. Note that if there was only one player, $a^{(1)}$ would represent an equilibrium. If this is not the case, let player 2 select an resource $r^{(2)}$ according to

$$r^{(2)} \in \arg \max_{r \in \mathcal{R}} f_r(2, \{a^{(1)}\}_r \cup \{2\}).$$

Denote the action profile $a^{(2)} = (r^{(1)}, r^{(2)}, \emptyset, \dots, \emptyset)$. If $r^{(1)} \neq r^{(2)}$, then by (9) and Condition 5.2 we know that

$$f_{r^{(1)}}(1, \{1\}) \geq f_{r^{(2)}}(1, \{1\}) \geq f_{r^{(2)}}(1, \{1, 2\}).$$

Therefore, player 1 can not improve his utility by altering his selection. If $r^{(1)} = r^{(2)} = r$, then by Condition 5.3, we know that for any resource $\bar{r} \in \mathcal{R}$,

$$\frac{f_{\bar{r}}(2, \{2\})}{f_{\bar{r}}(1, \{1\})} \geq \frac{f_r(2, \{1, 2\})}{f_r(1, \{1, 2\})}.$$

Using the above inequality, we can conclude that for any resource $\bar{r} \in \mathcal{R}$

$$f_r(2, \{1, 2\}) \geq f_{\bar{r}}(2, \{2\}) \Rightarrow f_r(1, \{1, 2\}) \geq f_{\bar{r}}(1, \{1\}).$$

Therefore, player 1 cannot improve his utility by altering his selection. Note that if there were only two players, $a^{(2)}$ would represent an equilibrium. Otherwise this argument could be repeated n times to construct an equilibrium.

□

It remains an open question as to whether Conditions 5.1 – 5.3 guarantee additional properties pertaining to the structure of the game besides existence of an equilibrium. For example, if the distribution rule satisfies Conditions 5.1 – 5.3, is the game a potential game or some variant?

5.2. Comparison with Existing Results

In Chen et al. (2008), the authors study cost sharing methodologies in a class of network formation games for a specific anonymous cost function. A network formation game is similar to a distributed welfare game where the difference lies in cost minimization versus welfare maximization; hence the results contained in that paper do not immediately translate to the framework of distributed welfare games. The authors results pertain to a specific anonymous resource specific cost function. Within their model, the authors show that if the distribution rule is (i) linear, which implies that the welfare is distributed similarly across resources, (ii) budget balance, and (iii) guarantees the existence of an equilibrium over all games then the distribution rule must correspond to a weighted Shapley value. Guaranteeing the existence of an equilibrium across all games means that if the distribution rule is fixed then an equilibrium is guaranteed to exist regardless of the number of resources or the structure of the players action set. To prove this results, the authors establish necessary pairwise conditions on player cost shares that are slightly stronger than the ones in Conditions 5.1 – 5.3. The authors make no reference as to whether their results also hold for alternative cost functions.

In Section 7.2, we demonstrate that our weaker pairwise conditions on player cost shares lead to the construction of a set of distribution rules that are budget balanced and guarantees an equilibrium in all games where players actions are singletons, i.e., $\mathcal{A}_i = \mathcal{R}$. Furthermore, the derived distribution rules do not correspond to a weighted Shapley value. This gap can potentially be a result of the following differences in the setup: (i) cost minimization versus welfare maximization, (ii) structure on action set, i.e., $\mathcal{A}_i = \mathcal{R}$ versus $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$, or (iii) structure of the welfare functions, i.e., anonymous versus non-anonymous. These differences highlight the fact that the limits of utility design for multi-agent systems is not well understood, hence it remains an open and interesting research direction.

6. Efficiency of Equilibria in Distributed Welfare Games

In addition to guaranteeing the existence of an equilibrium, it is important for a distribution rule to guarantee that the equilibria are efficient. In this section, we focus on bounding the efficiency of equilibria in distributed welfare games using the well known measures of *price of anarchy* (*PoA*) and *price of stability* (*PoS*) (Nissan et al. (2007)). In terms of distributed welfare games, the PoA gives a lower bound on the global welfare achieved by any equilibrium while the PoS gives a lower bound on the global welfare associated with the best equilibrium for any distributed welfare game. Specifically, let \mathcal{G} denote a set of distributed welfare games. For any particular game $G \in \mathcal{G}$ let $\mathcal{E}(G)$ denote the set of equilibria, $PoA(G)$ denote the price of anarchy, and $PoS(G)$ denote the price of stability for the game G where

$$PoA(G) := \min_{a^{ne} \in \mathcal{E}(G)} \frac{W(a^{ne})}{W(a^{opt})} \quad (10)$$

$$PoS(G) := \max_{a^{ne} \in \mathcal{E}(G)} \frac{W(a^{ne})}{W(a^{opt})}, \quad (11)$$

where $a^{opt} \in \arg \max_{a^* \in \mathcal{A}} W(a^*)$. We define the PoA and PoS for the set of distributed welfare games \mathcal{G} as

$$PoA(\mathcal{G}) := \inf_{G \in \mathcal{G}} PoA(G), \quad (12)$$

$$PoS(\mathcal{G}) := \inf_{G \in \mathcal{G}} PoS(G). \quad (13)$$

In general, the price of anarchy can be arbitrarily close to 0 in distributed welfare games. However, when the welfare function is submodular it is possible to attain a much better price of anarchy. We can interpret Theorem 3.4 in Vetta (2002) in the context of distributed welfare games to provide a fairly weak condition on the interaction between the welfare function W and the utility functions which guarantees that the price of anarchy is at least 1/2.

PROPOSITION 4. (Vetta (2002)) *Consider any distributed welfare game G . If for each resource $r \in \mathcal{R}$*

- (i) *the welfare function W_r is submodular,*
- (ii) *for each set of players $S \subseteq N$ and player $i \in S$, the distribution rule satisfies*

$$f_r(i, S) \geq W_r(S) - W_r(S \setminus \{i\}),$$

- (iii) *for each set of players $S \subseteq N$, the distribution rule satisfies*

$$\sum_{i \in S} f_r(i, S) \leq W_r(S)$$

then if an equilibrium exists the price of anarchy is 1/2.

To provide a basis for comparison, computing the optimal assignment for a general distributed welfare game with submodular welfare functions is NP-complete (Murphey (1999b)). Further, the best known approximation algorithms guarantee only to provide a solution that is within $1 - 1/e \approx 0.6321$ of the optimal (Feige and Vondrak (2006), Ageev and Sviridenko (2004), Ahuja et al. (2003)). Thus, the 1/2 price of anarchy in this scenario is comparable to the best centralized solution. It is important to note that these best known centralized approximations are in polynomial time whereas finding a Nash equilibrium is generally not polynomial time. However, recent results suggest that for this class of problems there are dynamics that get close to this 1/2 price of anarchy guarantee in polynomial time (Roughgarden (2009)).

While the generality of Proposition 4 is useful, the applicability is limited because it does not guarantee the existence of an equilibrium. Hence, its applicability depends on the results we have proven in Section 4.

COROLLARY 1. *Consider any distributed welfare game G where for each resource $r \in \mathcal{R}$ the welfare function W_r is submodular. If for each resources $r \in \mathcal{R}$*

- (i) the welfare function W_r is anonymous and the distribution rule f_r corresponds to the equally shared utility as in (3), or
 - (ii) the distribution rule f_r corresponds to the wonderful life utility as in (4), or
 - (iii) the distribution rule f_r corresponds to the (weighted) Shapley value as in (5) or (8),
- then an equilibrium exists and the price of anarchy is $1/2$.

The four distribution rules depicted in Corollary 1 all guarantee the existence of an equilibrium. Note that the wonderful life utility design satisfies condition (ii) with equality in addition to condition (iii) since the welfare function is submodular. Additionally, the Shapley and weighted Shapley values satisfies condition (iii) with equality and can easily be seen to satisfy condition (ii) when the welfare function is submodular. Finally, note that the price of anarchy is tight for these rules as illustrated in the following example.

EXAMPLE 2 (TIGHTNESS OF PRICE OF ANARCHY). Consider a distributed welfare game with player set $N = \{1, \dots, n\}$, resources $\mathcal{R} = \{r_1, \dots, r_n\}$, actions set $\mathcal{A}_i = \mathcal{R}$ for all players $i \in N$, and anonymous resource specific welfare functions of the form $W_{r_i}(S) = c_i$ for any player set $S \neq \emptyset$. Let $c_1 = 1$ and $c_2 = \dots = c_n = 1/n$. If the distribution rule is of the form (3), then an equilibrium is all characterized by all player choosing r_1 . The optimal allocation is all players choosing different resources. The efficiency of this situation is $n/(2n - 1)$ which goes to $1/2$ for large n . This example demonstrates that the equal share utility design and Shapley value utility design have a tight price of anarchy of $1/2$. Alternative examples can be constructed to show that the wonderful life utility and the weighted Shapley value utility also have a tight price of anarchy of $1/2$.

To this point we have focused exclusively on bounding the price of anarchy. Interestingly, when we focus on the price of stability there is a distinction between rules that are budget balanced and those that are not as highlighted in Table 6. If we require a budget balanced distribution rule, then both the price of anarchy and the price of stability is $1/2$. If budget balance is not a requirement, then it is possible to obtain a price of anarchy of $1/2$ and a price of stability is 1.

Distribution Rule	Budget Balanced	Price of Anarchy	Price of Stability
Equally Shared (anonymous)	yes	$1/2$	$1/2$
WLU	no	$1/2$	1
Shapley Value	yes	$1/2$	$1/2$

Table 2 Price of Anarchy and Price of Stability Comparison

6.1. Single Selection Distributed Welfare Games

In order to provide a tighter characterization of the price of anarchy we shift attention to single selection distributed welfare games. In this case, we can strengthen the results of Proposition 4 utilizing Conditions 5.1–5.3 which guarantee the existence of an equilibrium.

PROPOSITION 5. Consider any distributed welfare game G . If

- (i) for each resource $r \in \mathcal{R}$ the welfare function W_r is submodular,
 - (ii) for all players $i \in N$ the action sets satisfy $\mathcal{A}_i = \mathcal{R}$ and
 - (iii) the distribution rules $\{f_r\}$ are budget balanced and satisfy Conditions 5.1–5.3
- then an equilibrium exists and the price of anarchy is $1/2$.

Proof: The proof relies on showing that Conditions 5.1 – 5.3 combine to ensure that Condition (ii) of Proposition 4 is satisfied, i.e., for any resource $r \in \mathcal{R}$, set of players $S \subseteq N$, and player $i \in S$, we have

$$f_r(i, S) \geq W_r(S) - W_r(S \setminus \{i\}).$$

Let $i, j \in N$ be any two players where i is stronger than j . Let $S \subseteq N \setminus \{i, j\}$ be any player set. Condition 5.3 gives us that for any resource $r \in \mathcal{R}$,

$$f_r(j, S \cup \{j\}) \geq \frac{f_r(i, S \cup \{i\})}{f_r(i, S \cup \{i, j\})} f_r(j, S \cup \{i, j\}).$$

Since player i is stronger than player j , from Condition 5.2 we know that $f_r(i, S \cup \{i\}) \geq f_r(i, S \cup \{i, j\})$ which gives us

$$f_r(j, S \cup \{j\}) \geq f_r(j, S \cup \{i, j\}). \quad (14)$$

Therefore, Condition 5.2 holds for any players $i, j \in N$ and set of players $S \subseteq N \setminus \{i, j\}$.

Using the fact that the distribution rule is budget balanced and satisfies (14), we have

$$\begin{aligned} f_r(i, S) + W_r(S \setminus \{i\}) - W_r(S) &= f_r(i, S) + \sum_{j \in S \setminus \{i\}} f_r(j, S \setminus \{i\}) - \sum_{j \in S} f_r(j, S) \\ &= \sum_{j \in S \setminus \{i\}} f_r(j, S \setminus \{i\}) - f_r(j, S) \\ &\geq 0 \end{aligned}$$

Therefore, rearranging the above equation we have $f_r(i, S) \geq W_r(S) - W_r(S \setminus \{i\})$ which completes the proof.

□

6.2. Anonymous Distributed Welfare Games

Our bounds on the price of anarchy to this point have been independent of the number of players. In this section, we investigate the relationship between the price of anarchy and the number of players, albeit in the limited case where players are anonymous with regard to their impact on the global welfare. Furthermore, we analyze the price of anarchy when the number of players at the equilibrium and optimal allocations differ. Specifically, let $W(a^{\text{ne}}; n + \delta)$ be the total welfare garnered by an equilibrium consisting of $n + \delta$ players. Likewise, let $W(a^{\text{opt}}; n)$ be the total welfare garnered by an optimal allocation consisting of n players. While we allow variations in the number of players, the resources \mathcal{R} and their respective welfare functions $\{W_r\}$ remain fixed.

THEOREM 2. *Consider any distributed welfare game G . If*

- (i) *for each resource $r \in \mathcal{R}$ the welfare function W_r is anonymous and submodular,*
- (ii) *the action set of any two players $i, j \in N$ are identical, i.e., $\mathcal{A}_i = \mathcal{A}_j \subseteq 2^{\mathcal{R}}$,*
- (iii) *for any set of players $S \subseteq N$ and player $i \in N$, the distribution rule f_r assigns player i a share greater than or equal the marginal contribution,*

$$f_r(i, S) \geq W_r(S) - W_r(S \setminus \{i\}),$$

then if an equilibrium exists the relative price of anarchy satisfies

$$\frac{W(a^{\text{ne}}; n + \delta)}{W(a^{\text{opt}}; n)} \geq \frac{n + \delta}{2n + \delta - 1}.$$

Proof: We prove the result by bounding $W(a^{\text{opt}}; n)$ in terms of $W(a^{\text{ne}}; n + \delta)$. Rather than proving this theorem in terms of the distribution, we use the utility functions which are of the form

$$U_i(a_i, a_{-i}) = \sum_{r \in a_i} f_r(i, \{a\}_r).$$

Rewriting condition (iii) in terms of utility functions, we have that

$$U_i(a_i, a_{-i}) \geq W(a) - W(a_i^0, a_{-i}) \quad (15)$$

where $a_i^0 = \emptyset$.

First, notice that an upper bound on the $W(a^{\text{opt}}; n)$ is if one player in the optimal allocation can attain the entire welfare garnered at the equilibrium, $W(a^{\text{ne}}; n + \delta)$, and all other players attain $\min_{i \in N} U_i(a^{\text{ne}}; n + \delta)$ where $U_i(a^{\text{ne}}; n + \delta)$ represents the utility player i receives for the allocation a^{ne} consisting of $n + \delta$ players. To see that this upper bound holds, note first that (15) guarantees that each player's utility is an upper bound on the player's contribution to the global welfare. Further, by combining the definition of an equilibrium with the fact that the welfare function is submodular, we see that no additional player can attain a utility higher than $\min_{i \in N} U_i(a^{\text{ne}}; n + \delta)$ once $W(a^{\text{ne}}; n + \delta)$ is covered. Thus, we have

$$W(a^{\text{opt}}; n) \leq W(a^{\text{ne}}; n + \delta) + (n - 1) \min_{i \in N} U_i(a^{\text{ne}}; n + \delta).$$

Now, noting that

$$\min_{i \in N} U_i(a^{\text{ne}}; n + \delta) \leq \frac{W(a^{\text{ne}}; n + \delta)}{n + \delta}$$

gives

$$W(a^{\text{ne}}; n + \delta) + (n - 1) \min_{i \in N} U_i(a^{\text{ne}}; n + \delta) \leq W(a^{\text{ne}}; n + \delta) \left(1 + \frac{n - 1}{n + \delta} \right)$$

which easily gives the bound in the theorem

$$\begin{aligned} \frac{W(a^{\text{ne}}; n + \delta)}{W(a^{\text{opt}}; n)} &\geq \frac{1}{1 + \frac{n - 1}{n + \delta}} \\ &= \frac{n + \delta}{2n + \delta - 1}. \end{aligned}$$

□

Notice that Theorem 2 shows that the worst-case price of anarchy is increasing as the number of players increases and that as $n \rightarrow \infty$ the price of anarchy approaches $1/2$, which matches Proposition 4. Example 2 illustrates that this bound is tight by slightly modifying the coefficients to $c_2 = \dots = c_{n+\delta} = 1/(n + \delta)$. Furthermore, note that all the utility design methods previously studied, i.e., equally shared, wonderful life, and (weighted) Shapley value utility, satisfy the three conditions of Theorem 2. Hence, if the welfare function is submodular, then an equilibrium is guaranteed to exist and the bound on the relative price of anarchy holds. Lastly, note that the price of anarchy, $\delta = 0$, is bounded by

$$\frac{W(a^{\text{ne}}; n)}{W(a^{\text{opt}}; n)} \geq \frac{n}{2n - 1}.$$

7. Illustrative Examples

In this section we illustrate the developments in this paper on two well studied resource allocation problems: the graph coloring problem and the vehicle target assignment problem. Since both the graph coloring problem and the vehicle target assignment problem fit into the class of resource allocation problem studied in this paper, many of the methodologies established in this paper are directly applicable.

Our first goal in this section is to highlight that utility design for multi-agent does not need to be ad-hoc and application specific. To show this, we focus on the results in Panagopoulou and Spirakis (2008) where

the authors propose and analyze utility functions for the problem of graph coloring. We illustrate that the proposed design is equivalent to Shapley value utility design proposed in this paper. Utilizing the Shapley value utility design would have eliminated the need for having to prove existence of an equilibrium or a potential game structure. Furthermore, by recognizing their connection many of these results hold for my general setups.

Our second goal in this section is to highlight that many of the methodologies discussed in this paper provide strong efficiency guarantees in many settings. For example, in the vehicle target assignment problem all of the methodologies discussed in this paper guarantee that all equilibria are at least 50% efficient since the welfare functions are submodular. Furthermore, we demonstrate that the sufficient conditions established in Section 5 lead to the construction of equally desirable utility functions that are less demanding than either the marginal contribution or Shapley value utility functions. Lastly, we demonstrate how the structure of the specific welfare functions can be exploited to tighten the efficiency guarantees.

7.1. The Graph Coloring Problem

In the graph coloring problem there is a finite set of colors (or resources) denoted by \mathcal{C} and a graph represented by the tuple (N, E) where N is a finite number of nodes (or players) and $E \subseteq 2^{N \times N}$ is a set of directed edges on the graph G . Each node is allowed to choose any color, i.e., $\mathcal{A}_i = \mathcal{C}$ for all nodes $i \in N$. A color assignment is a tuple $a = (a_1, \dots, a_n)$ that associates a color with each node. We call a color assignment *valid* if $c_i \neq c_j$ for all nodes i, j such that $(e_i, e_j) \in E$. The goal of the graph coloring problem is to find a valid coloring assignment using the least number of possible colorings.

To formally describe the graph coloring problem in the context of distributed welfare games, we associate with each color $c \in \mathcal{C}$ a welfare function $W_c : 2^N \rightarrow \mathbb{R}$ where for any subset of nonconflicted players $S \subseteq N$, i.e. if $i, j \in S$ then $(e_i, e_j) \notin E$, we have

$$W_c(S) = \begin{cases} 0 & S = \emptyset \\ -1 & S \neq \emptyset \end{cases}$$

If S contains conflicted players, i.e., there exists players $i, j \in S$ such that $(e_i, e_j) \in E$, then we adopt the convention that $W_c(S) = -\infty$. The goal of the graph coloring problem is to find a coloring assignment a to maximize

$$W(a) = \sum_{c \in \mathcal{C}} W_c(\{a\}_c).$$

In Panagopoulou and Spirakis (2008), the authors model the graph coloring problem as a noncooperative game where each node is assigned a utility function of the form

$$U_i(a_i, a_{-i}) = \sum_{c \in a_i} |a|_c.$$

where a is any valid coloring assignment. Since this utility design was constructed specifically for the graph coloring problem the authors needed to prove results pertaining to existence and efficiency of equilibrium and the underlying potential game structure. It turns out that the proposed design is equivalent to assigning each player a utility in accordance with the Shapley value

$$\begin{aligned} U_i^{\text{SV}}(a_i, a_{-i}) &= \sum_{c \in a_i} f_c^{\text{SV}}(\{a\}_c) \\ &= \sum_{c \in a_i} -\frac{1}{|a|_c}. \end{aligned}$$

where a is any valid coloring assignment. By equivalence, we mean that for any assignment $a \in \mathcal{A}_i$, player $i \in N$, and alternative choice $a'_i \in \mathcal{A}_i$ we have

$$U_i(a_i, a_{-i}) - U_i(a'_i, a_{-i}) > 0 \Leftrightarrow U_i^{\text{SV}}(a_i, a_{-i}) - U_i^{\text{SV}}(a'_i, a_{-i}) > 0.$$

This equivalence implies that many of the results pertaining to equilibrium existence and the resulting potential game structure would be obtained for free as utilizing these methodologies eliminate the guess and check protocols commonly used in utility design. Lastly, since the Shapley value utility design provides guarantees irrespective of the structure of the players' action sets $\{\mathcal{A}_i\}$, many of the results hold immediately for more general setting. For example, the authors restrict attention to singleton strategies, i.e., $\mathcal{A}_i = C$, however, many of the results directly extend to more complex setting with non-singleton strategies, i.e. $\mathcal{A}_i \subseteq 2^C$, and non-symmetric action sets, i.e., $\mathcal{A}_i \neq \mathcal{A}_j$.

7.2. The Vehicle Target Assignment Problem

In the classical vehicle target assignment problem there is a finite set of targets (or resources) denoted by \mathcal{T} and each target $t \in \mathcal{T}$ has a relative worth $v_t \geq 0$. There are a finite number of vehicles (or players) denoted by N . The set of possible assignments for vehicle i is $\mathcal{A}_i \subseteq 2^{\mathcal{T}}$ and \mathcal{A} represents the set of joint assignments. Lastly, each vehicle $i \in N$ is parameterized with a success probability (depending on the context also referred to as detection, destroy, or completion probability) denoted by $p_i(t, a_i) \in [0, 1]$ which indicates the probability vehicle i will successfully eliminate target t given the assignment a_i . We assume that the success probabilities satisfy:

$$\begin{aligned} t \in a_i &\Leftrightarrow p_i(t, a_i) > 0, \\ t \notin a_i &\Leftrightarrow p_i(t, a_i) = 0. \end{aligned}$$

The goal of the vehicle target assignment problem is to find a joint assignment assign that maximizes the global welfare function

$$W(a) = \sum_{t \in \mathcal{T}} v_t \left(1 - \prod_{i \in \{a\}_t} [1 - p_i(t, a_i)] \right). \quad (16)$$

where $\left(1 - \prod_{i \in \{a\}_t} [1 - p_i(t, a_i)] \right)$ represents the joint probability of successfully eliminating target t .

To formulation of the vehicle target assignment is identical to several alternative resource allocation problems ranging from sensor coverage to task allocation. Furthermore, computing the optimal assignment for this class of problems is an NP-hard combinatorial optimization problem Murphey (1999b). Resultantly, research has traditionally centered around developing heuristic methods to quickly obtain near optimal assignment, where the degree of suboptimality is dependent on the structure of the global objective, e.g., Ahuja et al. (2003).

7.2.1. A Game Theoretic Formulation: Note that the vehicles target assignment problem is precisely a resource allocation problem with a separable welfare function where the welfare function for any target $t \in \mathcal{T}$ and any set of vehicles $S \subseteq N$ is

$$W_t(S) = v_t \left(1 - \prod_{i \in S} [1 - p_i(t, a_i)] \right). \quad (17)$$

Therefore, we can appeal to the utility design methodologies developed for distributed welfare game to construct local utility for the vehicles such that the resulting game has a host of desirable properties including existence and efficiency of equilibrium. One possible design choice is the marginal contribution distribution rule which takes on the form

$$\begin{aligned} U_i(a_i, a_{-i}) &= \sum_{t \in a_i} f_t^{\text{MC}}(i, \{a\}_t) \\ &= \sum_{t \in a_i} v_t \left(p_i(t, a_i) \prod_{j \in \{a\}_t \setminus \{i\}} (1 - p_j(t, a_j)) \right). \end{aligned} \quad (18)$$

An alternative design choice is the weighted Shapley value distribution rule where for a given set of player weights $w \in R_+^n$ the utility takes on the form

$$\begin{aligned} U_i(a_i, a_{-i}) &= \sum_{t \in a_i} f_t^{\text{WSV}}(i, \{a\}_t) \\ &= \sum_{t \in a_i} v_t \left(\sum_{S \subseteq \{a\}_t: i \in S} \frac{w_i}{\sum_{j \in S} w_j} \left(\sum_{R \subseteq T} (-1)^{|T|-|R|} \left(1 - \prod_{j \in R} (1 - p_j(t, a_j)) \right) \right) \right) \end{aligned} \quad (19)$$

where the Shapley value can be attained with $w_i = 1$ for all $i \in N$. Both (18) and (19) guarantee the existence of an equilibrium. Further, since the welfare function is submodular, both rules yield a price of anarchy of $1/2$. Further, the price of stability of the wonderful life design is 1 and the price of stability of the weighted Shapley design is $1/2$. There are many issues that may limit the practical applicability of the wonderful life and (weighted) Shapley designs. The most important issues are that the wonderful life design is not budget balanced and that, in general, computing the (weighted) Shapley values is not tractable.

7.2.2. Extending Equal Share to Weighted Share: Are we bound to using the (weighted) Shapley value if the goal is to design utility functions that are budget balanced and guarantee the existence of an equilibrium? Or is it possible to extend the equal share design in (3) to a weighted share design to satisfy these properties? By weighted share, we mean assigning each player i a strength coefficient $s_i > 0$ where the distribution rule at any target t takes on the form

$$f_t(i, S) = \frac{s_i}{\sum_{j \in S} s_j} W_t(S) \quad (20)$$

for any subset of vehicles $S \subseteq N$. For example, a logical choice in the case of the vehicles target assignment problem would be for $s_i = p_i$.

If we restrict attention to the case where each vehicle can select only a single target, i.e., $\mathcal{A}_i = \mathcal{T}$, and each vehicle $i \in N$ has an invariant success probability $p_i > 0$, i.e., if $a_i = t$ then $p_i(t, a_i) = p_i$ otherwise $p_i(t, a_i) = 0$, then this weighted share distribution rule always guarantees the existence of an equilibrium as the following theorem demonstrates.

THEOREM 3. Consider any distributed welfare game $G = \{N, \mathcal{T}, \{\mathcal{A}_i\}, \{W_t\}, \{f_t\}\}$. If

- (i) the welfare function for each target $t \in \mathcal{T}$ takes on the form in (17),
- (ii) the action set of each vehicles is $\mathcal{A}_i = \mathcal{T}$,
- (iii) the distribution rule for each target $t \in \mathcal{T}$ for any set of vehicles $S \subseteq N$ and player $i \in S$ is

$$f_t(i, S) = \frac{p_i}{\sum_{j \in S} p_j} W_t(S)$$

then the resulting utility functions are budget balanced, an equilibrium exists, and the price of anarchy is $1/2$.

Proof: We prove this result by verifying Conditions 5.1-5.3. First, Condition 5.1 is satisfied trivially since for any target $t \in \mathcal{T}$, set of vehicles $S \subseteq N$ and vehicles $i, j \in S$ we have

$$\frac{f_t(i, S)}{f_t(j, S)} = \frac{p_i}{p_j}.$$

Verifying Condition 5.2 requires showing that for any target $t \in \mathcal{T}$, set of vehicles $S \subseteq N$ and vehicles $i \in S$ and $j \notin S$ where $p_i > p_j$, we have

$$\frac{p_i}{p_S} W_t(S) \geq \frac{p_i}{p_j + p_S} W_t(S \cup \{j\})$$

where $p_S = \sum_{k \in S} p_k$. Using the fact that $W_t(S \cup \{j\}) = W_t(S) + p_j(v_r - W_t(S))$ and rearranging the above expression we need to show that

$$p_S(v_t - W_t(S)) \leq W_t(S).$$

Working with the lefthand side of the above expression, we have

$$p_S(v_t - W_t(S)) = \sum_{i \in S} p_i(v_t - W_t(S)) \tag{21}$$

$$\leq \sum_{i \in S} p_i(v_t - W_t(S \setminus \{i\})) \tag{22}$$

$$\leq \sum_{i \in S} \frac{p_i}{p_S} W_t(S) \tag{23}$$

$$= W_t(S) \tag{24}$$

where the first and step follow from submodularity of W_t .

Verifying Condition 5.3 requires showing that for any target $t \in \mathcal{T}$, set of vehicles $S \subseteq N$, and vehicles $i, j \notin S$ where $p_i > p_j$, we have

$$\frac{f_t(j, S \cup \{j\})}{f_t(j, S \cup \{i, j\})} \geq \frac{f_t(i, S \cup \{j\})}{f_t(i, S \cup \{i, j\})}$$

which is equivalent to showing that

$$(p_i + p_S)W_t(S \cup \{j\}) \geq (p_j + p_S)W_t(S \cup \{i\}).$$

Using the fact that $W_t(S \cup \{i\}) = W_t(S) + p_i(v_t - W_t(S))$ for any player $i \notin S$ this is equivalent to showing that

$$p_S(v_t - W_t(S)) \leq W_t(S)$$

which is true from the previous analysis in (21)–(24).

□

7.2.3. Remarks on Theorem 3: A few notes are in order as to the meaning of the resulting utility design in Theorem 3. First, note that the utility design set forth in Theorem 3 is budget balanced and guarantees the existence of an equilibrium regardless of the game setup, i.e., the number of vehicles, their respective success probabilities, or the number of targets provided that the action sets satisfy $\mathcal{A}_i = \mathcal{T}$. Furthermore, this utility has an information dependency similar to that of equal share utility design in (3).

This utility design has several interesting properties. First of all, setting the strength as $s_i = p_i$ for each vehicles is not the unique choice of strengths that guarantee equilibrium existence when using the weighted share distribution rule in (20). In fact, for any fixed $k \in [0, 1]$, letting the strength of each vehicle $i \in N$ be defined by the solutions to the equation

$$p_i = (1 - k)s_i + k \frac{s_i}{1 + s_i}$$

guarantees the same properties as in Theorem 3 with a similar proof that we omit for brevity. The importance of this is that there are family of strength coefficient that guarantee equilibrium existence; hence, this strength based design scheme is not a razor edge phenomenon. Understanding when such strength based decompositions is possible is fundamentally important to understanding utility design in multiagent systems. From a design perspective, having a complete understanding of this space of admissible utility function is important as it allows the designer to optimize over this class.

Lastly, this result in some ways contradicts the results in Chen et al. (2008) which suggest that in general using the weighted Shapley value is the only distribution rules that guarantees the existence of an equilibrium and is budget balanced. It is straightforward to show that the utility design in Theorem 3 does not corresponds to a weighted Shapley value for any fixed weight vector $w \in R_+^n$. It is fundamentally important to understand the root of the discrepancy which could depend on the differences in the setup: (i) cost minimization versus welfare maximization, (ii) structure on action set, i.e., $\mathcal{A}_i = \mathcal{R}$ versus $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$, or (iii) structure of the welfare functions, i.e., anonymous versus non-anonymous.

7.2.4. Efficiency Improvement: In this section we demonstrate that frequently the price of anarchy guarantees can be strengthened by exploiting the structure of the welfare functions. To simplify our analysis, we focus purely on the anonymous single selection case where each vehicle has the same success probability p . In this setting, we can directly appeal to Theorem 2 to show that the relative price of anarchy is

$$\frac{W(a^{\text{ne}}; n + \delta)}{W(a^{\text{opt}}; n)} \geq \frac{n + \delta}{2n + \delta - 1}.$$

We can strengthen these bounds further to attain the following bound, which illustrates the impact of the success probability on the price of anarchy.

THEOREM 4. Consider any distributed welfare game $G = \{N, \mathcal{T}, \{\mathcal{A}_i\}_{i \in N}, \{W_t\}_{t \in \mathcal{T}}, \{f_t\}_{t \in \mathcal{T}}\}$. If

- (i) the welfare function for each target $t \in \mathcal{T}$ is anonymous and takes on the form in (17),
- (ii) the action set of each vehicles is $\mathcal{A}_i = \mathcal{T}$,
- (iii) each vehicle $i \in N$ has the same detection probability $p_i = p$ where $p \in [0, 1]$,
- (vi) for any set of vehicles $S \subseteq N$ and player $i \in S$ the distribution rule for each target $t \in \mathcal{T}$ satisfies

$$f_t(i, S) = \frac{1}{|S|} W_t(S)$$

then the resulting utility functions are budget balanced, an equilibrium exists, and the price of anarchy is bounded by

$$\frac{W(a^{\text{ne}})}{W(a^{\text{opt}})} \geq \left(\frac{a^*}{n} + \frac{1 - (1-p)^{n-a^*}}{1 - (1-p)^n} \right)^{-1}$$

$$\text{where } a^* = \begin{cases} n - 1, & p = 1; \\ n - \frac{\log\left(n \frac{\log(1/(1-p))}{1 - (1-p)^n}\right)}{\log(1/(1-p))}, & p < 1. \end{cases}$$

Figure 1 illustrate the price of anarchy. Note that this bound is a decreasing function of the vehicles success probability which is intuitive.

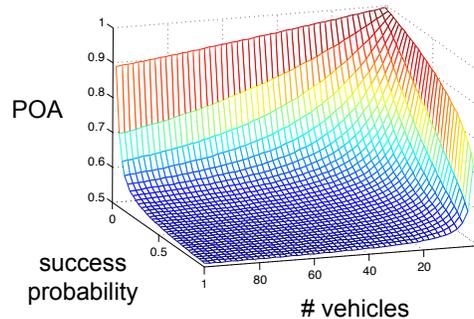


Figure 1 Price of Anarchy in Single Selection Anonymous Vehicle Target Assignment Game

7.2.5. Proof of Theorem 4: Theorem 4 follows from the following sequence of lemmas. Throughout we let n represent the number of vehicles and p the invariant success probability.

LEMMA 1. *The price of anarchy is bounded by*

$$\left(\max_{x+y \leq n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^{x+y}} \right\} \right)^{-1}$$

where the maximum is taken over integer x, y .

Proof: We describe the optimal assignment in terms of the Nash assignment. From the Nash conditions and the fact that each vehicle's utility is greater than or equal to the marginal contribution, we have that $|a^{\text{ne}}|_t > 0 \Rightarrow |a^{\text{opt}}|_t > 0$. We can bound the welfare of the optimal as

$$\begin{aligned} W(a^{\text{opt}}) &\leq \sum_{t \in \mathcal{T}: |a^{\text{ne}}|_t > 0} \max_{m_t \in [0, |a^{\text{ne}}|_t - 1]} \left\{ m_t \frac{W_t(|a^{\text{ne}}|_t)}{|a^{\text{ne}}|_t} + W_t(|a^{\text{ne}}|_t - m_t) \right\} \\ &= \sum_{t \in \mathcal{T}: |a^{\text{ne}}|_t > 0} \max_{m_t \in [0, |a^{\text{ne}}|_t - 1]} \left\{ \frac{m_t}{|a^{\text{ne}}|_t} + \frac{(1 - (1-p)^{|a^{\text{ne}}|_t - m_t})}{(1 - (1-p)^{|a^{\text{ne}}|_t})} \right\} W_t(|a^{\text{ne}}|_t) \end{aligned}$$

which follows from the fact that each vehicle's utility is greater than or equal to the vehicle's marginal contribution to the welfare. Letting $x_t = m_t$ and $y_t = |a^{\text{ne}}|_t - m_t$, we have

$$\begin{aligned} W(a^{\text{opt}}) &\leq \sum_{t \in \mathcal{T}: |a^{\text{ne}}|_t > 0} \max_{x_t + y_t = |a^{\text{ne}}|_t, x_t \geq 0, y_t \geq 1} \left\{ \frac{x_t}{x_t + y_t} + \frac{1 - (1-p)^{y_t}}{1 - (1-p)^{x_t}} \right\} W_t(|a^{\text{ne}}|_t) \\ &\leq \sum_{t \in \mathcal{T}: |a^{\text{ne}}|_t > 0} \max_{x+y \leq n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^x} \right\} W_t(|a^{\text{ne}}|_t) \\ &= \max_{x+y \leq n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^x} \right\} \sum_{t \in \mathcal{T}: |a^{\text{ne}}|_t > 0} W_t(|a^{\text{ne}}|_t) \\ &= \max_{x+y \leq n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^x} \right\} W(a^{\text{ne}}) \end{aligned}$$

which completes the proof.

□

To obtain a more explicit form of the price of anarchy, we first relax the constraints followed by characterized the maximal x, y .

LEMMA 2.

$$\max_{x+y \leq n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^{x+y}} \right\} \leq \max_{x+y=n, x \geq 0, y \geq 1} \left\{ \frac{x}{x+y} + \frac{1 - (1-p)^y}{1 - (1-p)^{x+y}} \right\}$$

where the LHS is taken over integer x, y and the RHS is taken over real-valued x, y .

Proof: We start by relaxing the integer optimization to include real-valued x, y . Next, suppose, that $x_m + y_m = m < n$ are the maximizers under the constraint that $x + y = m$. We will show that $x_n = nx_m/m$, $y_n = ny_m/m$ lead to a larger value than x_m, y_m . Combining this with the observation that $x_n + y_n = n$ then completes the proof.

$$\frac{x_n}{x_n + y_n} + \frac{1 - (1-p)^{y_n}}{1 - (1-p)^{x_n + y_n}} = \frac{x_m}{x_m + y_m} + \frac{1 - (1-p)^{(n/m)y_m}}{1 - (1-p)^{(n/m)(x_m + y_m)}}$$

Now, it is enough to show that the right hand side is increasing in (n/m) since $n > m$, which follows from simple calculations.

□

Now, we know that $y = n - x$. So, we need only calculate x .

LEMMA 3.

$$x^* = \arg \max_{0 \leq x \leq n-1} \left\{ \frac{x}{n} + \frac{1 - (1-p)^{n-x}}{1 - (1-p)^n} \right\} = \begin{cases} n-1, & p = 1; \\ n - \frac{\log\left(\frac{n \log(1/(1-p))}{1 - (1-p)^n}\right)}{\log(1/(1-p))}, & p < 1. \end{cases}$$

Proof: For the case of $p = 1$, the result is immediate. In the case when $p \neq 1$, we evaluate the maximizer by simply differentiating. Differentiating with respect to x gives:

$$\frac{1}{n} - \frac{(1-p)^{n-x} \log(1/(1-p))}{1 - (1-p)^n}$$

Setting the derivative equal to zero, then gives

$$(1-p)^{n-x} = \frac{n \log(1/(1-p))}{1 - (1-p)^n}$$

Solving for x , we obtain

$$x = n - \frac{\log\left(\frac{n \log(1/(1-p))}{1 - (1-p)^n}\right)}{\log(1/(1-p))}$$

which completes the proof.

□

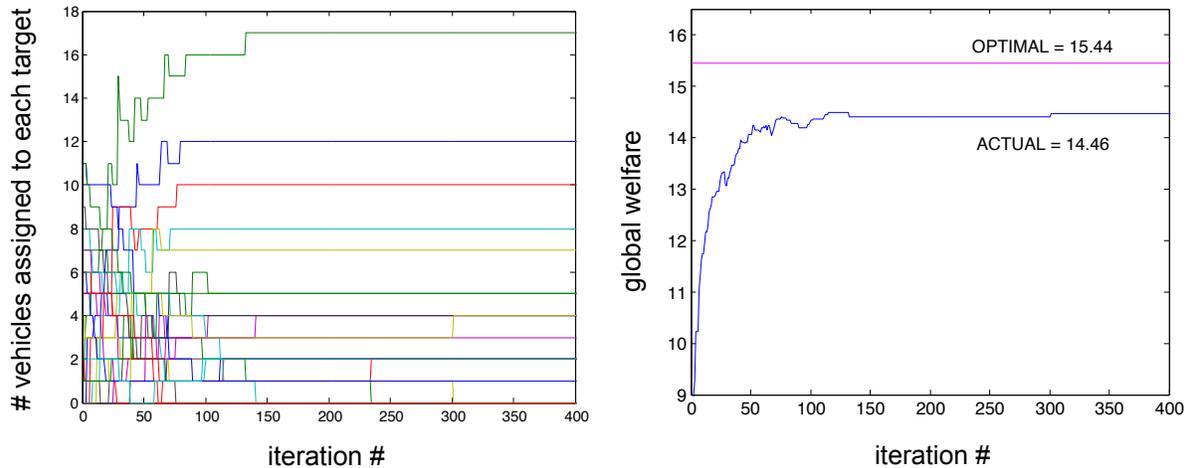
7.3. Simulation Experiments

To this point, we have explored equilibrium behavior in the game theoretic formulation of the vehicle target assignment problem. The question that remains is how can the autonomous vehicles reach an equilibrium in a distributed fashion. While not focusing on this question in detail, we illustrate the applicability of the theory of learning in games as a local control mechanism for coordinating behavior in such systems.

We consider a vehicle target assignment problem with 25 targets $\mathcal{T} = \{t_1, \dots, t_{25}\}$ and 100 vehicles each capable of assigning itself to just a single target, i.e., $\mathcal{A}_i = \mathcal{T}$. Each vehicle has an invariant success probability $p = 0.25$. The reward for each target, v_t , is randomly assigned from a uniform distribution; two targets according to $U[0, 6]$, four targets according to $U[0, 3]$, and the remaining according to $U[0, 1]$. Each vehicle is assigned an equally shared utility design (3).

There is a large body of literature analyzing distributed learning algorithms in potential games (Young (1998), Monderer and Shapley (1996a), Marden et al. (2009, 2007b,a)). We apply *fading memory joint strategy fictitious play with inertia*, which guarantees convergence to an equilibrium in any potential game while maintaining computational tractability even in large-scale games (Marden et al. (2009)). Fading memory joint strategy fictitious play with inertia can be described as follows:

1. **Initialization:** Each vehicle i is assigned a perceived utility $V_i^{a_i}(1) \in \mathbb{R}$ for each action $a_i \in \mathcal{A}_i$. One can think of $V_i^{a_i}(1)$ as vehicles i 's initial belief about the utility he would receive for playing action a_i in the ensuing time step.
2. **Action Selection:** At time $t \geq 1$, each vehicle i plays the following strategy:
 - $a_i(t) \in \arg \max_{a_i \in \mathcal{A}_i} V_i^{a_i}(t)$ with probability $(1 - \epsilon)$,
 - $a_i(t) = a_i(t-1)$ with probability (ϵ) ,



(a) Evolution of number of vehicles assigned to each target during mission

(b) Evolution of welfare function during mission

Figure 2 Simulation results for vehicle target assignment problem.

where $\epsilon \in (0, 1)$ is referred to as the vehicle's inertia, i.e., a probabilistic reluctance to changing strategies.

3. **Belief Propagation:** Each vehicle i updates his beliefs as

$$V_i^{a_i}(t+1) = (1 - \lambda)V_i^{a_i}(t) + \lambda U_i(a_i, a_{-i}(t)), \quad \forall a_i \in \mathcal{A}_i,$$

where $\lambda \in (0, 1]$ is the discount factor.

4. Return to Step 2 and repeat.

It is worth noting that fading memory joint strategy fictitious play with inertia only requires each vehicle to observe the number of vehicles assigned to each target. The identity of the observed vehicles is unimportant. If all vehicles adhere to the prescribed learning rule, then the action profile, $a(t)$, will converge almost surely to an equilibrium.

We use the following discount factor and inertia: $\lambda = 0.5$ and $\epsilon = 0.02$. Figure 2(a) illustrates the evolution of the number of vehicles assigned to each target. The identity of the targets is unimportant as the key observation is that behavior settles down at an equilibrium. Figure 2(b) illustrates the evolution of the global welfare in addition to the efficiency gap between the equilibrium and the optimal. From Theorem 4, we know that the price of anarchy must be greater than 0.541. Our simulation illustrates that Theorem 4 provides a very conservative estimate of the efficiency since the price of anarchy we observe is 0.936.

8. Concluding Remarks

This paper represents the first comprehensive effort at understanding utility design for multi-agent systems. To formally study utility design we introduce a class of games that we refer to as distributed welfare games. We demonstrate that cost sharing methodologies are beneficial for utility design; however, such methodologies do not typically guarantee all desirable properties of interest. While this paper addresses the applicability of cost sharing methodologies for utility design, significant future work entails understanding the limits of utility design. In particular, is it impossible to satisfy all performance metrics simultaneously? If so, how can this be accomplished? Formally understanding the limits of utility design for noncooperative games is of fundamental importance to understand how game theory should evolve to meet the challenges inherent within the control of multi-agent systems.

9. Acknowledgements

The authors would like to acknowledge the support of Sherwin Doroudi, Kenneth McKell, Microsoft Research, and both the Social and Information Sciences Laboratory and the Lee Center for Advanced Networking at Caltech.

References

- Ageev, A., M. Sviridenko. 2004. Pipe rounding: a new method for constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization* **8** 307–328.
- Ahuja, R. K., A. Kumar, K. Jha, J. B. Orlin. 2003. Exact and heuristic methods for the weapon-target assignment problem. [Http://ssrn.com/abstract=489802](http://ssrn.com/abstract=489802).
- Akella, A., R. Karp, C. Papadimitriou, S. Seshan, S. Shenker. 2002. Selfish behavior and stability of the internet: A game-theoretic analysis of tcp. *Proceedings of SIGCOMM*.
- Arslan, G., J. R. Marden, J. S. Shamma. 2007. Autonomous vehicle-target assignment: a game theoretical formulation. *ASME Journal of Dynamic Systems, Measurement and Control* **129** 584–596.
- Chen, H-L., T. Roughgarden, G. Valiant. 2008. Designing networks with good equilibria. *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. 854–863.
- Conitzer, V., T. Sandholm. 2004. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. *Proceedings of AAI*.
- E. Campos-Nañez, C. Li, A. Garcia. 2008. A game-theoretic approach to efficient power management in sensor networks. *Operations Research* **56**(3) 552–561.
- Feige, U., J. Vondrak. 2006. Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 667–676.
- Gairing, M. 2009. Covering games: Approximation through non-cooperation. *Proceedings of the Fifth Workshop on Internet and Network Economics (WINE)*.
- Garcia, A., D. Reaume, R. L. Smith. 2000. Fictitious play for finding system optimal routings in dynamic traffic networks. *Transportation Research Part B* **34**(2) 147–156.
- Goemans, M., L. Li, V. S. Mirrokni, M. Thottan. 2004. Market sharing games applied to content distribution in ad-hoc networks. *Symposium on Mobile Ad Hoc Networking and Computing (MOBIHOC)*.
- Gopalakrishnan, R., J. R. Marden, A. Wierman. 2010. An architectural view of game theoretic control. *Proceedings of ACM HotMetrics*.
- Haeringer, G. 2006. A new weight scheme for the shapley value. *Mathematical Social Sciences* **52**(1) 88–98.
- Hart, S., A. Mas-Colell. 1989. Potential, value, and consistency. *Econometrica* **57**(3) 589–614.
- Johari, R., J. N. Tsitsiklis. 2004. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research* **29**(3) 407–435.
- Kaumann, B., F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, C. Diot. 2007. Measurement based self organization of interfering 802.11 wireless access networks. *Proceedings of INFOCOM*.
- Komali, R. S., A. B. MacKenzie. 2007. Distributed topology control in ad-hoc networks: A game theoretic perspective. *Proceedings of IEEE Consumer Communication and Network Conference*.
- Krause, A., C. Guestrin. 2007. Near-optimal observation selection using submodular functions. *Proc. of Conf. on Artificial Intelligence*.
- Li, W., C. G. Cassandras. 2005. Sensor networks and cooperative control. *European Journal of Control* To appear.
- Marden, J. R., G. Arslan, J. S. Shamma. 2007a. Connections between cooperative control and potential games illustrated on the consensus problem. *Proceedings of the 2007 European Control Conference (ECC '07)*.
- Marden, J. R., G. Arslan, J. S. Shamma. 2007b. Regret based dynamics: Convergence in weakly acyclic games. *Proceedings of the 2007 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Honolulu, Hawaii.
- Marden, J. R., G. Arslan, J. S. Shamma. 2009. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control* **54** 208–220.

- Marden, J. R., M. Effros. 2009. The price of selfishness in network coding. Under submission.
- Marden, J. R., A. Wierman. 2008. Distributed welfare games. Under submission.
- Marden, J. R., H. P. Young, G. Arslan, J. S. Shamma. 2007c. Payoff based dynamics for multi-player weakly acyclic games. *Proceedings of the 46th IEEE Conference on Decision and Control*.
- Mhatre, V., K. Papagiannaki, F. Baccelli. 2007. Interference mitigation through power control in high density 802.11. *Proceedings of INFOCOM*.
- Monderer, D., L. Shapley. 1996a. Fictitious play property for games with identical interests. *Games and Economic Theory* **68** 258–265.
- Monderer, D., L. Shapley. 1996b. Potential games. *Games and Economic Behavior* **14** 124–143.
- Murphey, R. A. 1999a. Target-based weapon target assignment problems. P. M. Pardalos, L. S. Pitsoulis, eds., *Nonlinear Assignment Problems: Algorithms and Applications*. Kluwer Academic, Alexandria, Virginia, 39–53.
- Murphey, R.A. 1999b. Target-based weapon target assignment problems. P.M. Pardalos, L.S. Pitsoulis, eds., *Nonlinear Assignment Problems: Algorithms and Applications*. Kluwer Academic Publishers, 39–53.
- Nissan, N., T. Roughgarden, E. Tardos, V. V. Vazirani. 2007. *Algorithmic game theory*. Cambridge University Press, New York, NY, USA.
- Panagopoulou, P., P. Spirakis. 2008. A game theoretic approach for efficient graph coloring. S.-H. Hong, N. Nagamochi, T. Fukunaga, eds., *Lecture notes in computer science*. Springer-Verlag, 183–195.
- Rosenthal, R. W. 1973. The network equilibrium problem in integers. *Networks* **3**(1) 53–59.
- Roughgarden, T. 2005. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, MA, USA.
- Roughgarden, T. 2009. Intrinsic robustness of the price of anarchy. *Proceedings of STOC*.
- Sandholm, W. 2002. Evolutionary implementation and congestion pricing. *Review of Economic Studies* **69**(3) 667–689.
- Shapley, L.S. 1953. A value for n -person games. H. W. Kuhn, A. W. Tucker, eds., *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)*. Princeton University Press, Princeton, NJ, 307–317.
- Srivastava, V., J. Neel, A. MacKenzie, J. Hicks, L.A. DaSilva, J.H. Reed, R. Gilles. 2005. Using game theory to analyze wireless ad hoc networks. *IEEE Communications Surveys and Tutorials* To appear.
- Vetta, A. 2002. Nash equilibria in competitive societies with applications to facility location, traffic routing, and auctions. *Proc. of Symp. on Fdns. of Comp. Sci.* 416–425.
- Wolpert, D., K. Tumor. 1999. An overview of collective intelligence. J. M. Bradshaw, ed., *Handbook of Agent Technology*. AAAI Press/MIT Press.
- Young, H. P. 1994. *Equity*. Princeton University Press, Princeton, NJ.
- Young, H. P. 1998. *Individual Strategy and Social Structure*. Princeton University Press, Princeton, NJ.
- Zou, Y., K. Chakrabarty. 2004. Uncertainty-award and coverage-oriented deployment for sensor networks. *Journal of Parallel and Distributing Computing* **64**(7) 788–798.