

Distributed Welfare Games

Jason R. Marden

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

Adam Wierman

Computing and Mathematical Sciences, California Institute of Technology, Pasadena, California 91125,
adamw@caltech.edu

Game-theoretic tools are becoming a popular design choice for distributed resource allocation algorithms. A central component of this design choice is the assignment of utility functions to the individual agents. The goal is to assign each agent an admissible utility function such that the resulting game possesses a host of desirable properties, including scalability, tractability, and existence and efficiency of pure Nash equilibria. In this paper we formally study this question of utility design on a class of games termed distributed welfare games. 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 commonly studied classes of resource allocation problems: “coverage” problems and “coloring” problems.

Subject classifications: resource allocation; game theory; distributed control.

Area of review: Revenue Management.

History: Received July 2010; revisions received August 2011, August 2012; accepted September 2012. Published online in *Articles in Advance* February 8, 2013.

1. Introduction

The central challenge in any resource allocation problem is to allocate a fixed number of resources such that their utilization is optimized. Successfully dealing with this challenge is essential for ensuring reliable and efficient performance in a host of applications ranging from management of transportation networks to the routing of information through the internet. Traditionally, researchers have focused on developing centralized algorithms to determine efficient allocations. However, in many modern applications, these centralized algorithms are neither applicable or desirable.

One concrete example highlighting the need for distributed resource allocation is the “sensor coverage problem,” 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 (see, e.g., Li and Cassandras 2005, Martinez et al. 2007). A centralized algorithm for sensor allocation requires that a central authority maintain complete knowledge of the environment and communicate directly with each sensor during the entire mission. Both requirements might be unrealistic in large and/or hostile environments. Similar issues arise in many computer network resource allocation problems, with examples ranging from wireless access point assignment (Kauffmann et al. 2007) to wireless power management (Campos-Náñez et al. 2008, Li and Cassandras 2005). There are also many examples outside of computer systems that suffer from similar issues. For example, in transportation systems a global planner does not have the authority to assign drivers to roads; rather, a global planner must entice drivers

appropriately to settle on a desirable allocation, e.g., using tolls to minimize aggregate congestion (Sandholm 2002).

The need for distributed resource allocation has led to a surge of research aimed at understanding the possibility of decentralizing (localizing) decisions in resource allocation problems. This is an extremely diverse literature where protocols are designed using a wide variety of tools, e.g., distributed optimization (Mishra et al. 2006, Villegas et al. 2008), distributed control (Li et al. 2005, Parag et al. 2010), physics-inspired control (Kauffmann et al. 2007, Mhatre et al. 2007), and game-theoretic control (Alpcan et al. 2009, Zou and Chakrabarty 2004, Srivastava et al. 2005, Campos-Náñez et al. 2008).

In this paper we focus on game-theoretic control, which is a promising new approach for distributed resource allocation. A game-theoretic approach to distributed resource allocation requires two distinct design steps. First, a system designer must model the interaction framework of the decision-making entities as a strategic form game. This involves specifying the decision makers, their respective choices, and a local utility function for each agent. Second, a system designer must specify a local behavioral or learning rule for each agent that specifies how an individual agent processes available information to formulate a decision. The overarching goal is to complete the two design steps, referred to as *utility design* and *learning design*, respectively, to ensure that the emergent global behavior is desirable (Arslan et al. 2007, Marden et al. 2009, Gopalakrishnan et al. 2011).

There are wide-ranging advantages to the game-theoretic approach, including robustness to failures and

environmental disturbances, minimal communication requirements, improved scalability, and real-time adaptation (Arslan et al. 2007). Accordingly, game-theoretic resource allocation designs are increasingly popular in a variety of wireless and sensor network applications, e.g., channel access control in wireless networks (Altman et al. 2004, Kauffmann et al. 2007), coverage problems in sensor networks (Cassandras and Li 2005), and power control in both (Altman and Altman 2003, Campos-Náñez et al. 2008, Falomari et al. 1999). A comprehensive survey of applications can be found in Altman et al. (2006). However, nearly all these designs are highly application-specific, with both the utility and learning designs crafted carefully for the specific setting.

Our Contributions

The goal of this paper is to establish a general framework for studying utility design. To that end, we introduce a class of games termed *distributed welfare games*, which represents a game theoretic model for resource allocation problems with separable objective/welfare functions. Here, we focus on the design of *local* agent utility functions only where local means that an agent's utility function is able to depend only on the resources selected, the welfare at each resource, and the other agents that selected the same resources. Consequently, the utility design question can be viewed as a welfare sharing problem where an agent's utility is defined as some fraction of the welfare garnered at each of the agent's selected resources. Therefore, designing local utility functions is equivalent to defining a distribution rule that depicts how the welfare garnered from each resource is distributed to the players at that particular resource.

The first set of results (§3) illustrates 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 (pure) 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-balanced utility functions and tractability.

Our second set of results seeks to overcome the limitations of the Shapley value by establishing tractable budget-balanced distribution rules, (§4). We identify three sufficient conditions on distribution rules, see Conditions 4.1–4.3, 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 illustrate these developments in §6.

Our third set of results pertains to the efficiency of the resulting Nash equilibria. When restricting attention to submodular welfare functions, which is a common attribute to many resource allocation problems (Vetta 2002, Krause and Guestrin 2007), we identify distribution rules that guarantee that all resulting Nash equilibria obtain a welfare within 50% of the welfare associated with the optimal assignment. 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. 2004). Surprisingly, this comparison demonstrates that the inefficiency resulting from localizing decisions in resource allocation problems is relatively small when the welfare functions are submodular. Furthermore, we demonstrate that these 50% guarantees can be significantly sharpened in many settings. Lastly, in §6 we provide an illustration of the theoretical contributions contained in this paper on two important classes of resource allocation problems: coverage problems and coloring problems.

2. Model Overview: Resource Allocation Problems

We consider a 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. Let $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$ denote the set of joint action profiles. 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 1999), content distribution (Goemans et al. 2004), graph coloring (Panagopoulou and Spirakis 2008), and network coding (Marden and Effros 2012), 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).

We now introduce the class of *distributed welfare games* as the game-theoretic model for the class of resource allocation problems defined above. We consider a finite strategic-form game where each player $i \in N$ has a finite action sets \mathcal{A}_i and a utility function $U_i: \mathcal{A} \rightarrow \mathbb{R}$. 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 agent i for any joint action profile $a \in \mathcal{A}$ is defined as

$$U_i(a) = \sum_{r \in a_i} f_r(i, \{a\}_r), \quad (1)$$

where $f_r: N \times 2^N \rightarrow \mathbb{R}$ 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 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 (1). First, this structure imposes a natural notion of locality as each agent's utility function depends only on the resources the agent selected and the other agents that selected the same resources. Second, 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 because we can remove the dependence between utility design, or equivalently distribution rule design, and the structure of the players' action sets.

3. 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 our context can be viewed as the inverse of a cost-sharing problem (Young 1994, Shapley 1953, Hart and Mas-Colell 1989). However, while cost-sharing methodologies can be effective as distribution rules, we also discuss several issues that limit their applicability.

3.1. Performance Criteria

Before proceeding with the discussion of distribution rules, we first formally define our performance criteria. 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 $r \in \mathcal{R}$ depend *only* 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.

Existence and efficiency of pure Nash equilibrium. A well-known equilibrium concept that emerges in strategic form games is that of a *pure Nash equilibrium*, which we will refer to simply as an equilibrium. An action profile $a^* \in \mathcal{A}$ is called a *equilibrium* if for all players $i \in N$,

$$U_i(a_i^*, a_{-i}^*) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^*), \quad (2)$$

where we adopt the convention that a_{-i} denotes 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 sometimes write a profile a of actions as (a_i, a_{-i}) . Similarly, we may write $U_i(a)$ as $U_i(a_i, a_{-i})$. A pure Nash equilibrium, as defined in (2), represents a scenario for which no player has a unilateral incentive to deviate. 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. The class of potential games (Monderer and Shapley 1996) imposes a restriction on the agents' 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 $a'_i, a''_i \in \mathcal{A}_i$,

$$U_i(a'_i, a_{-i}) - U_i(a''_i, a_{-i}) = \phi(a'_i, a_{-i}) - \phi(a''_i, a_{-i}). \quad (3)$$

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 an equilibrium, hence every potential game possesses at least one such equilibrium. 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 $S \subseteq N$, $\sum_{i \in S} f_r(i, S) = W_r(S)$. Recent results demonstrate that having budget-balanced utility functions (or in this case distribution rules), or some fraction thereof, is important for providing desirable efficiency guarantees (Vetta 2002, Roughgarden 2009, Gairing 2009). Furthermore, having budget-balanced distribution rules is also important for the control (or influence) of social systems 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) and content distribution (Goemans et al. 2004).

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 follows:

- *High:* The distribution to player i given the player set S is conditioned on information regarding the welfare associated with all player set $T \subseteq S$, i.e., $f_r(i, S; \{W_r(T)\}_{T \subseteq S})$.
- *Medium:* The distribution to player i given the player set S is conditioned on information regarding the welfare associated with the player sets S and $S \setminus \{i\}$, i.e., $f_r(i, S; \{W_r(S), W_r(S \setminus \{i\})\})$.
- *Low:* The distribution to player i given the player set S is conditioned on information regarding the welfare associated with the set S , i.e., $f_r(i, S; W_r(S))$.

3.2. Distribution Rules

In this section we will introduce four distribution rules motivated by methodologies from the cost-sharing literature (Young 1994). These rules are summarized in Table 1.

3.2.1. Equally Shared. The equally shared distribution rule takes on the following form: for any resource $r \in \mathcal{R}$, player set $S \subseteq N$, and player $i \in N$,

$$f_r^{\text{ES}}(i, S) = \frac{W_r(S)}{|S|}. \quad (4)$$

It is straightforward to verify that this rule is budget-balanced and has a low information dependency. However, in general such a design does not guarantee the existence of an equilibrium (Arslan et al. 2007). However, if players are *anonymous* with regard to their impact on the welfare functions, 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 (4) guarantee the existence of an equilibrium. We refer to such welfare functions as anonymous.

PROPOSITION 1. *If G is a distributed welfare game with anonymous welfare functions and the equally share distribution rule in (4), then an equilibrium is guaranteed to exist.*

PROOF. 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 1996), with the following specification: (a) resources \mathcal{R} ; (b) cost functions of the form

$c_r(k) = (W_r(k))/k$, $k > 0$, where k is the number of players utilizing resource r ; and (c) utility functions of the form $U_i(a) = \sum_{r \in a_i} c_r(|a|_r)$. Because any congestion game is a potential game, this completes the proof. \square

3.2.2. Marginal Contribution. The marginal contribution distribution rule takes on the following form: for any resource $r \in \mathcal{R}$, player set $S \subseteq N$, and player $i \in N$,

$$f_r^{\text{MC}}(i, S) = W_r(S) - W_r(S \setminus \{i\}). \quad (5)$$

Note that the marginal contribution distribution rule requires a medium informational dependency as each agent is required to compute the player's marginal contribution to the welfare. 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 (5) results in a potential game with potential function W ; hence any action profile that maximizes the global welfare is an equilibrium. However, other equilibria might also exist. Furthermore, the marginal contribution distribution rule may distribute more (or less) welfare than is gathered; hence, it is not budget-balanced.

3.2.3. Shapley Value. The Shapley value distribution rule (Shapley 1953, Hart and Mas-Colell 1989, Haeringer 2006) takes on the following form: for any resource $r \in \mathcal{R}$, player set $S \subseteq N$, and player $i \in N$

$$f_r^{\text{SV}}(i, S) = \sum_{T \subseteq S \setminus \{i\}} \frac{|T|!(|S| - |T| - 1)!}{|S|!} (W_r(T \cup \{i\}) - W_r(T)). \quad (6)$$

Utilizing the Shapley value as in (6) requires a high informational dependency; however, it rectifies the budget-balanced problems associated with the marginal contribution distribution rule as shown in the following proposition.

PROPOSITION 2. *If G is a distributed welfare game with the Shapley value distribution rule in (6) then the players' utility functions are budget balanced and the game is a potential game with potential function*

$$\phi^{\text{SV}}(a) := \sum_{r \in \mathcal{R}} \sum_{S \subseteq \{a\}_r} \frac{1}{|S|} \left(\sum_{T \subseteq S} (-1)^{|S| - |T|} W_r(T) \right). \quad (7)$$

PROOF. Budget-balanced utility functions follow directly from Hart and Mas-Colell (1989) and Haeringer (2006). We prove the second part of this proposition using the potential function derived in Hart and Mas-Colell (1989). First, we express the Shapley value distribution rule in (6) as a weighted sum of unanimity games (Hart and Mas-Colell 1989, Haeringer 2006), which takes on the form

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

Let $\alpha_r(T) := \sum_{R \subseteq T} (-1)^{|T|-|R|} W_r(R)$ and $\phi_r^{\text{SV}}(a) := \sum_{T \subseteq \{a\}_r} (\alpha_r(T)) / |T|$ be the resource-specific potential function (Hart and Mas-Colell 1989). Furthermore, let $a \in \mathcal{A}$ be any allocation and $a_i^0 = \emptyset$ be the null action for player i . Using (6), 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 $a_{-i} \in \mathcal{A}_{-i}$ we have

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

which completes the proof. \square

There are two limitations of the Shapley value utility design that could 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 (6) or (8), where computation of the Shapley value requires a weighted summation over all subsets of players, of which there might be exponentially many. However, it should be noted that this computational cost is lessened dramatically for special classes of welfare functions, e.g., Conitzer and Sandholm (2004). For example, if players are anonymous, then the Shapley value is equivalent to the equal share distribution rule in (4).

3.2.4. 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. The weighted Shapley value distribution rule (Shapley 1953, Hart and Mas-Colell 1989, Haeringer 2006) takes on the following form: for any resource $r \in \mathcal{R}$, player set $S \subseteq N$, and player $i \in N$,

$$f_r^{\text{WSV}}(i, S) := \sum_{T \subseteq S: i \in T} \frac{w_i}{\sum_{j \in T} w_j} \left(\sum_{R \subseteq T} (-1)^{|T|-|R|} W_r(R) \right). \quad (9)$$

Note that the Shapley value distribution rule in (6) is recaptured if $w_i = 1$ for all players $i \in N$. We state the following proposition without proof to avoid redundancy.

PROPOSITION 3. *If G is a distributed welfare game with the Shapley value distribution rule in (9), then the players' utility functions are budget balanced and the game 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 (7). However, as with the Shapley value, the potential function can be computed recursively and is of the form $\phi^{\text{WSV}}(a) := \sum_{r \in \mathcal{R}} \phi_r^{\text{WSV}}(\{a\}_r)$, where (Hart and Mas-Colell 1989)

$$\begin{aligned} \phi_r^{\text{WSV}}(\emptyset) &= 0, \\ \phi_r^{\text{WSV}}(S) &= \frac{1}{\sum_{i \in S} w_i} \left[W_r(S) + \sum_{i \in S} w_i \phi_r^{\text{WSV}}(S \setminus \{i\}) \right], \\ &\quad \forall S \subseteq N, S \neq \emptyset. \end{aligned}$$

3.3. Comparison of Distribution Rules

The following table summarizes the features of the four distribution rules introduced above. Note that there is an inherent trade-off between the desirable features of a distribution rule and the computational and informational requirement needed to obtain such a rule.

4. Single Selection Distributed Welfare Games

The results of the previous section suggest that deriving 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 trade-off 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 allowed to select only a single resource, $\mathcal{A}_i = \mathcal{R}$, as opposed to multiple resources, $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. Although this setting is simplified, there are still a wide variety of resource allocation problems that typically are modeled as single selection resource allocation problems, e.g., task allocation.

4.1. 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 4.1. 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

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

Table 1. Summary of distribution rules for distributed welfare games.

Distribution rule	Existence of equilibrium	Potential game	Budget balanced	Tractable	Informational requirement
Equally shared (anonymous)	Yes	Yes	Yes	Yes	Low
Equally shared WLU	No	No	Yes	Yes	Low
Shapley	Yes	Yes	No	Yes	Medium
Weighted Shapley	Yes	Yes	Yes	No	High

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 4.2. 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 4.3. 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. If G is a distributed welfare game where the action sets satisfy $\mathcal{A}_i = \mathcal{R}$ for all agents $i \in N$ and the distribution rule satisfies Conditions 4.1–4.3, 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 4.1. 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 a resource $r^{(1)}$ according to

$$r^{(1)} \in \arg \max_{r \in \mathcal{R}} f_r(1, \{1\}). \quad (10)$$

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 a 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 (10) and Condition 4.2 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 4.3, we know that for any resource $\bar{r} \in \mathcal{R}$,

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

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

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

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. \square

It remains an open question as to whether Conditions 4.1–4.3 guarantee additional properties pertaining to the structure of the game besides existence of an equilibrium.

4.2. Comparison with Existing Results

The related work in Chen et al. (2008) studies 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 focus on a specific anonymous cost function and prove that a distribution rule is budget balanced and guarantees the existence of an equilibrium for any game if and only if the distribution rule can be represented by a weighted Shapley value. To prove this result, the authors establish necessary and sufficient pairwise conditions on player distributed shares that are slightly stronger than the ones in Conditions 4.1–4.3. The authors make no claim as to whether their results also hold for alternative cost functions.

In §6.2, 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 nonanonymous.

5. Efficiency of Equilibria in Distributed Welfare Games

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) (Nisan 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

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

$$PoS(G) := \max_{a^{ne} \in \mathcal{E}(G)} \frac{W(a^{ne})}{W(a^{opt})}, \quad (12)$$

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 $PoA(\mathcal{G}) := \inf_{G \in \mathcal{G}} PoA(G)$ and $PoS(\mathcal{G}) := \inf_{G \in \mathcal{G}} PoS(G)$.

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). *If G is a distributed welfare game, where 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 greater than or equal to $1/2$.

To provide a basis for comparison, computing the optimal assignment for a general distributed welfare game with submodular welfare functions is NP-complete (Murphey 1999). Furthermore, 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. 2004). 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 §3.

COROLLARY 1. *If G is a distributed welfare game where for each resource $r \in \mathcal{R}$:*

- (i) *the distribution rule f_r corresponds to the marginal contribution as in (5), or*
 - (ii) *the distribution rule f_r corresponds to the (weighted) Shapley value as in (6) or (9),*
- then an equilibrium exists and the price of anarchy is greater than or equal to $1/2$.*

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) because 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. The following examples highlights a specific distributed welfare game meeting the conditions set forth in Corollary 1, which yields a price of anarchy of $1/2$; hence, the only way to attain a price of anarchy $> 1/2$ is to impose additional structure on the game environment, which we will explore in the ensuing section.

EXAMPLE 1 (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 (4), then an equilibrium is all characterized by all players 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 price of anarchy (and price of stability) 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 2. When considering the class of distributed welfare games with submodular welfare functions, the price of anarchy and the price of stability is $1/2$ for any budget balanced distribution rule. 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.

Table 2. PoA and PoS comparison over all distributed welfare games with submodular welfare functions.

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

5.1. Single Selection Distributed Welfare Games

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 4.1–4.3, which guarantee the existence of an equilibrium.

PROPOSITION 5. *If G is a distributed welfare game where:*

- (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 4.1–4.3,*
- then an equilibrium exists and the price of anarchy is 1/2.*

PROOF. The proof relies on showing that Conditions 4.1–4.3 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 4.3 gives us that for any resource $r \in \mathcal{R}$,

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

Because player i is stronger than player j , from Condition 4.2 we know that $f_r(i, S \cup \{i\}) \geq f_r(i, S \cup \{i, j\})$, which gives us

$$f_r(j, S \cup \{j\}) \geq f_r(j, S \cup \{i, j\}). \quad (13)$$

Therefore, Condition 4.2 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 (13), 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, we have $f_r(i, S) \geq W_r(S) - W_r(S \setminus \{i\})$, which completes the proof. \square

5.2. 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 differs. Specifically, let $W(a^{ne}; n + \delta)$ be the total welfare garnered by an equilibrium consisting of $n + \delta$ players, $\delta \in \mathbb{Z}^+$. Likewise, let $W(a^{opt}; n)$ 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. *If G is a distributed welfare game where*

- (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 satisfies*

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

then if an equilibrium exists, the relative price of anarchy satisfies

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

PROOF. We prove the result by bounding $W(a^{opt}; n)$ in terms of $W(a^{ne}; n + \delta)$. Rather than proving this theorem in terms of the distribution rules, we use the utility functions, which are of the form $U_i(a_i, a_{-i}) = \sum_{r \in \mathcal{A}_i} f_r(i, \{a\}_r)$. Rewriting condition (iii) in terms of utility functions, letting $a_i^0 = \emptyset$ we have that

$$U_i(a_i, a_{-i}) \geq W(a) - W(a_i^0, a_{-i}). \quad (14)$$

First, notice that an upper bound on the $W(a^{opt}; n)$ is if one player in the optimal allocation can attain the entire welfare garnered at the equilibrium, $W(a^{ne}; n + \delta)$, and all other players attain $\min_{i \in N} U_i(a^{ne}; n + \delta)$, where $U_i(a^{ne}; n + \delta)$ 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 (14) guarantees that each player's utility is an upper bound on the player's contribution to the global welfare. Furthermore, 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 $\min_{i \in N} U_i(a^{ne}; n + \delta)$ once $W(a^{ne}; n + \delta)$ is covered. Thus, we have

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

Now, noting that

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

gives

$$\begin{aligned} & W(a^{nc}; n + \delta) + (n - 1) \min_{i \in N} U_i(a^{nc}; n + \delta) \\ & \leq W(a^{nc}; n + \delta) \left(1 + \frac{n - 1}{n + \delta} \right), \end{aligned}$$

which easily gives the bound in the theorem

$$\frac{W(a^{nc}; n + \delta)}{W(a^{opt}; n)} \geq \frac{1}{1 + (n - 1)/(n + \delta)} = \frac{n + \delta}{2n + \delta - 1}. \quad \square$$

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 1 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

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

6. Illustrative Examples

To illustrate the relevance of the results discussed to this point, we now focus on two broad classes of resource allocation problems: coverage problems and coloring problems. Our first goal in this section is to highlight that utility design for distributed resource allocation does not need to be ad-hoc and application specific, it can often follow immediately from the general framework presented here. To illustrate 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, which highlights that utilizing the Shapley value utility design would have eliminated the need for having to prove existence of an equilibrium or a potential game structure.

Our second goal in this section is to highlight that many of the methodologies discussed in this paper provide strong efficiency guarantees in broad settings. For example, in coverage problems, all 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 §4 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.

6.1. Graph Coloring Problems

A graph coloring problem is defined as follows. 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.

Graph coloring problems play a prominent role in several class of resource allocation problems ranging from distributed caching (Chun et al. 2004) to spectrum allocation in cognitive radio networks (Moscibroda and Wattenhofer 2005, Schneider and Wattenhofer 2009), and the game theoretic techniques have recently been suggested as a useful approach for developing distributed protocols for such problems (Panagopoulou and Spirakis 2008, Chatzigiannakis et al. 2010).

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. Because 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

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

where a is any valid coloring assignment. By equivalent, 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

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

This equivalence implies that many of the results pertaining to equilibrium existence and the resulting potential game structure would be obtained for free because utilizing these methodologies eliminates the guess and check protocols commonly used in utility design. Lastly, because 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 a more general setting.

6.2. Coverage Problems

In a coverage problem there is a finite set of resources denoted by \mathcal{T} , and each location/resource $t \in \mathcal{T}$ has a relative value $v_t \geq 0$. There are a finite number of agents denoted by N . The set of possible assignments for agent i is $\mathcal{A}_i \subseteq 2^{\mathcal{T}}$ and \mathcal{A} represents the set of joint assignments. Lastly, each agent $i \in N$ is parameterized with a success probability denoted by $p_i(t, a_i) \in [0, 1]$, which indicates the probability that agent i will successfully “cover” resource t given the assignment a_i . We assume that the success probabilities satisfy

$$t \in a_i \Leftrightarrow p_i(t, a_i) > 0.$$

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

$$W(a) = \sum_{i \in \mathcal{T}: \{a_i\} \neq \emptyset} v_i \left(1 - \prod_{i \in \{a_i\}} [1 - p_i(t, a_i)] \right), \quad (15)$$

where $(1 - \prod_{i \in \{a_i\}} [1 - p_i(t, a_i)])$ represents the expected value that resource t is covered by the assignment a . Note that computing the optimal assignment for this class of problems is an NP-hard combinatorial optimization problem (Murphey 1999) and, 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. (2004).

To view coverage problems in the context of distributed welfare games, we simply note that they are resource allocation problems with a separable welfare function, where the welfare function for any location/resource $t \in \mathcal{T}$ and any set of agents $S \subseteq N$ is

$$W_t(S) = v_t \left(1 - \prod_{i \in S} [1 - p_i(t, a_i)] \right). \quad (16)$$

Therefore, we can appeal to the utility design methodologies developed for distributed welfare game to construct

local utility for the agents 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

$$U_i(a_i, a_{-i}) = \sum_{t \in a_i} f_t^{MC}(i, \{a\}_t), \\ = \sum_{t \in a_i} v_t \left(p_i(t, a_i) \prod_{j \in \{a_i\} \setminus \{i\}} (1 - p_j(t, a_j)) \right). \quad (17)$$

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

$$U_i(a_i, a_{-i}) = \sum_{t \in a_i} f_t^{WSV}(i, \{a\}_t) \\ = \sum_{t \in a_i} v_t \left(\sum_{S \subseteq \{a_i\}: i \in S} \frac{w_i}{\sum_{j \in S} w_j} \left(\sum_{R \subseteq T} (-1)^{|T|-|R|} \right. \right. \\ \left. \left. \cdot \left(1 - \prod_{j \in R} (1 - p_j(t, a_j)) \right) \right) \right), \quad (18)$$

where the Shapley value can be attained with $w_i = 1$ for all $i \in N$. Both (17) and (18) guarantee the existence of an equilibrium. Furthermore, because the welfare function is submodular, both rules yield a price of anarchy of $1/2$. Furthermore, the price of stability of the wonderful life design is 1 and the price of stability of the weighted Shapley design is $1/2$.

6.2.1. Other Distribution Rules. We have just seen that cost sharing methodologies yield useful distribution rules for coverage problems. However, a question that remains is whether we are 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. The following theorem identifies one alternative distribution rule, albeit for the case of singleton strategies.

THEOREM 3. *If $G = \{N, \mathcal{T}, \{\mathcal{A}_i\}, \{W_t\}, \{f_t\}\}$ is a distributed welfare game, where:*

- (i) *the welfare function for each location $t \in \mathcal{T}$ takes on the form in (16),*
- (ii) *the action set of each agent is $\mathcal{A}_i = \mathcal{T}$,*
- (iii) *the distribution rule for each location $t \in \mathcal{T}$ for any set of agents $S \subseteq N$ and player $i \in S$ is*

$$f_t(i, S) = \frac{P_i}{\sum_{j \in S} P_j} W_t(S), \quad (19)$$

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

PROOF. We prove this result by verifying Conditions 4.1–4.3. First, Condition 4.1 is satisfied trivially

because for any location $t \in \mathcal{T}$, set of agents $S \subseteq N$ and agents $i, j \in S$ we have

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

Verifying Condition 4.2 requires showing that for any location $t \in \mathcal{T}$, set of agents $S \subseteq N$ and agents $i \in S$ 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

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

Working with the left-hand side of the above expression, we have

$$\begin{aligned} p_S(v_r - W_t(S)) &= \sum_{i \in S} p_i(v_r - W_t(S)) \leq \sum_{i \in S} p_i(v_r - W_t(S \setminus \{i\})) \\ &\leq \sum_{i \in S} \frac{p_i}{p_S} W_t(S) = W_t(S), \end{aligned} \quad (20)$$

where the first and step follow from submodularity of W_t .

Verifying Condition 4.3 requires showing that for any location $t \in \mathcal{T}$, set of agents $S \subseteq N$, and agents $i, j \notin S$ where $p_i > p_j$, we have

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

Using the fact that $W_t(S \cup \{i\}) = W_t(S) + p_i(v_r - W_t(S))$ for any player $i \notin S$ this is equivalent to showing that $p_S(v_r - W_t(S)) \leq W_t(S)$, which is true from the previous analysis in (20). \square

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 agents, their respective coverage probabilities, or the number of locations, 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 (4).

This utility design has several interesting properties. First of all, one can view the distribution rule in (19) as a “weighted” share distribution rule of the general form

$$f_r(i, S) = \frac{s_i}{\sum_{j \in S} s_j} W_r(S),$$

where s_i is the strength (or weight) of player i . While setting $s_i = p_i$ is the obvious choice, it turns out that letting

the strength of each agent $i \in N$ be defined by the solutions to the equation

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

where $1 \geq k \geq 0$ also provides the same guarantees as in Theorem 3 with a similar proof that we omit for brevity. The importance of this is that there are family-of-strength coefficients that guarantee equilibrium existence; hence, this weighted share distribution rule is not a razor-edge phenomenon. Understanding when weighted shared distribution rules are possible is fundamentally important to understanding utility design. 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 sense contradicts the results in Chen et al. (2008) that suggest that the weighted Shapley value is the *only* distribution rule 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 correspond to a weighted Shapley value for any fixed weights $w_i \in R^+$. It is fundamentally important to understand whether the root of the discrepancy arises from (i) cost minimization versus welfare maximization, (ii) structure of the 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 nonanonymous.

6.2.2. Efficiency Improvement. So far, we have approached coverage problems using only the general results provided earlier in the paper. However, it is important to note that the price of anarchy guarantees provided earlier can frequently be strengthened by exploiting the structure (or curvature) of the welfare function for the specific application of interest. Here, we illustrate such a tightening in the context of coverage problems. To simplify our analysis, we focus purely on the anonymous single selection case where each agent 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


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

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

THEOREM 4. *If $G = \{N, \mathcal{T}, \{\mathcal{A}_i\}, \{W_t\}, \{f_t\}\}$ is a distributed welfare game where*

- (i) *the welfare function for each location $t \in \mathcal{T}$ is anonymous and takes on the form in (16),*
- (ii) *the action set of each agent is $\mathcal{A}_i = \mathcal{T}$,*
- (iii) *each agent $i \in N$ has the same detection probability $p_i = p$ where $p \in [0, 1]$,*

Figure 1. Price of anarchy in single selection, anonymous, coverage games.



(iv) for any set of agents $S \subseteq N$ and agent $i \in S$ the distribution rule for each location $t \in \mathcal{T}$ satisfies

$$f_i(i, S) = \frac{1}{|S|} W_i(S),$$

then the utility functions are budget balanced, an equilibrium exists, and the price of anarchy satisfies

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

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

Figure 1 illustrates the price of anarchy. Note that this bound is a decreasing function of the agents success probability, which is intuitive.

6.2.3. Proof of Theorem 4. Theorem 4 follows from the following sequence of lemmas. Throughout, we let n represent the number of agents and p represent 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 agent’s utility is greater than or equal to the marginal contribution, we have that $|a^{nc}|_t > 0 \Rightarrow |a^{opt}|_t > 0$. We can bound the welfare of the optimal as

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

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

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

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

which completes the proof. \square

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

LEMMA 2.

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

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.

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

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

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

LEMMA 3.

$$x^* = \operatorname{argmax}_{0 \leq x \leq n-1} \left\{ \frac{x}{n} + \frac{1 - (1-p)^{n-x}}{1 - (1-p)^n} \right\}$$

$$= \begin{cases} n-1, & p = 1; \\ n - \frac{\log((n \log(1/(1-p)))/(1 - (1-p)^n))}{\log(1/(1-p))}, & p < 1. \end{cases}$$

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

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

which completes the proof. \square

7. Concluding Remarks

This paper represents a comprehensive effort at understanding utility design for distributed resource allocation. 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. Some recent progress on this topic can be found in Gopalakrishnan et al. (2011) and Marden and Wierman (2013), but many open questions remain.

Acknowledgments

The authors thank the reviewers and editors of the paper for their valuable feedback, which improved the paper significantly. This work was supported by Air Force Office of Scientific Research [Grant FA9550-12-1-0359], Office of Naval Research [Grant N00014-12-1-0643], and the National Science Foundation [Grant CNS 0846025].

References

Ageev A, Sviridenko M (2004) Pipe rounding: A new method for constructing algorithms with proven performance guarantee. *J. Combin. Optim.* 8:307–328.

- Ahuja RK, Kumar A, Jha K, Orlin JB (2004) Exact and heuristic methods for the weapon-target assignment problem. MIT Sloan Working Paper 4464-03, Massachusetts Institute of Technology, Cambridge, MA.
- Alpcan T, Pavel L, Stefanovic N (2009) A control theoretic approach to noncooperative game design. *IEEE Conf. Decision and Control* (IEEE, Piscataway, NJ), 8575–8580.
- Altman E, Altman Z (2003) S -modular games and power control in wireless networks. *Trans. Automatic Control* 48(5):839–842.
- Altman E, El-Azouzi R, Jiménez T (2004) Slotted ALOHA as a game with partial information. *Comput. Networks* 45(6):701–713.
- Altman E, Boulogne T, El-Azouzi R, Jiménez T, Wynter L (2006) A survey on networking games in telecommunications. *Comput. Oper. Res.* 33(2):286–311.
- Arslan G, Marden JR, Shamma JS (2007) Autonomous vehicle-target assignment: A game theoretical formulation. *ASME J. Dynam. Systems, Measurement and Control* 129:584–596.
- Campos-Náñez E, Garcia A, Li C (2008) A game-theoretic approach to efficient power management in sensor networks. *Oper. Res.* 56(3):552–561.
- Cassandras CG, Li W (2005) Sensor networks and cooperative control. *Eur. J. Control* 11(4–5):436–463.
- Chatzigiannakis I, Koninis C, Panagopoulou PN, Spirakis PG (2010) Distributed game-theoretic vertex coloring. *Proc. Internat. Conf. Principles of Distributed Systems*.
- Chen H-L, Roughgarden T, Valiant G (2008) Designing networks with good equilibria. *Proc. Nineteenth Annual ACM-SIAM Sympos. Discrete Algorithms* (ACM, New York), 854–863.
- Chun B-C, Chaudhuri K, Wee H, Barreno M, Papadimitriou CH, Kubiawicz J (2004) Selfish caching in distributed systems: A game-theoretic analysis. *Proc. ACM Principles of Distributed Comput.* (ACM, New York).
- Conitzer V, Sandholm T (2004) Computing Shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. *Proc. 19th AAAI Conf. Artificial Intelligence* (AAAI, Palo Alto, CA).
- Falomari D, Mandayam N, Goodman D, Shah V (1999) A new framework for power control in wireless data networks: Games, utility, and pricing. *Wireless Multimedia Network Tech.* 289–310.
- Feige U, Vondrak J (2006) Approximation algorithms for allocation problems: Improving the factor of $1 - 1/e$. *47th Annual IEEE Sympos. Foundations Comput. Sci. (Sympos. Foundations Comput. Sci. '06)* (IEEE, Piscataway, NJ), 667–676.
- Gairing M (2009) Covering games: Approximation through non-cooperation. *Proc. Fifth Workshop on Internet and Network Econom. (WINE)*.
- Goemans M, Li L, Mirrokni VS, Thottan M (2004) Market sharing games applied to content distribution in ad-hoc networks. *Sympos. Mobile Ad Hoc Networking Comput. (MOBIHOC)*.
- Gopalakrishnan R, Marden JR, Wierman A (2011) An architectural view of game theoretic control. *ACM SIGMETRICS Performance Eval. Rev.* 38(3):31–36.
- Haeringer G (2006) A new weight scheme for the Shapley value. *Math. Soc. Sci.* 52(1):88–98.
- Hart S, Mas-Colell A (1989) Potential, value, and consistency. *Econometrica* 57(3):589–614.
- Kauffmann B, Baccelli F, Chaintreau A, Mhatre V, Papagiannaki K, Diot C (2007) Measurement-based self organization of interfering 802.11 wireless access networks. *IEEE INFOCOM* (IEEE, Piscataway, NJ), 1451–1459.
- Krause A, Guestrin C (2007) Near-optimal observation selection using submodular functions. *Proc. Conf. Artificial Intelligence* (AAAI, Palo Alto, CA).
- Li LE, Halpern JY, Bahl P, Wang Y-M, Wattenhofer R (2005) A cone-based distributed topology-control algorithm for wireless multi-hop networks. *Trans. Networking* 13(1):147–159.
- Li W, Cassandras CG (2005) Sensor networks and cooperative control. *IEEE Conf. Decision and Control* (IEEE, Piscataway, NJ), 4237–4238.
- Marden JR, Effros M (2012) The price of selfishness in network coding. *IEEE Trans. Inform. Theory* 58(4):2349–2361.

- Marden JR, Wierman A (2013) Overcoming the limitations of utility design for multiagent systems. *IEEE Trans. Automatic Control*. Forthcoming.
- Marden JR, Arslan G, Shamma JS (2009) Connections between cooperative control and potential games. *IEEE Trans. Systems, Man and Cybernetics. Part B: Cybernetics* 39:1393–1407.
- Martinez S, Cortes J, Bullo F (2007) Motion coordination with distributed information. *Control Systems Magazine* 27(4):75–88.
- Mhatre V, Papagiannaki K, Baccelli F (2007) Interference mitigation through power control in high density 802.11 WLANs. *IEEE INFOCOM* (IEEE, Piscataway, NJ), 535–543.
- Mishra A, Shrivastava V, Agrawal D, Banerjee S, Ganguly S (2006) Distributed channel management in uncoordinated wireless environments. *Conf. Mobile Comput. Networking* 170–181.
- Monderer D, Shapley L (1996) Potential games. *Games Econom. Behav.* 14:124–143.
- Moscibroda T, Wattenhofer R (2005) Coloring unstructured radio networks. *ACM Sympos. Parallelism in Algorithms and Architectures* (ACM, New York).
- Murphey RA (1999) Target-based weapon target assignment problems. Pardalos PM, Pitsoulis LS, eds. *Nonlinear Assignment Problems: Algorithms and Applications* (Kluwer Academic Publishers, Boston), 39–53.
- Nisan N, Roughgarden T, Tardos E, Vazirani VV (2007) *Algorithmic Game Theory* (Cambridge University Press, New York).
- Panagopoulou P, Spirakis P (2008) A game theoretic approach for efficient graph coloring. Hong S-H, Nagamochi N, Fukunaga T, eds. *Lecture Notes in Computer Science* (Springer-Verlag, Berlin), 183–195.
- Parag P, Shakkottai S, Chamberland J-F (2010) Value-aware resource allocation for service guarantees in networks. *IEEE INFOCOM* (IEEE, Piscataway, NJ).
- Rosenthal RW (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).
- Roughgarden T (2009) Intrinsic robustness of the price of anarchy. *Proc. 41st Annual ACM Sympos. Theory of Comput.* (ACM, New York).
- Sandholm W (2002) Evolutionary implementation and congestion pricing. *Rev. Econom. Stud.* 69(3):667–689.
- Schneider J, Wattenhofer R (2009) Coloring unstructured wireless multi-hop networks. *Proc. ACM Principles of Distributed Comput.* (ACM, New York).
- Shapley LS (1953) A value for n -person games. Kuhn HW, Tucker AW, eds. *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)* (Princeton University Press, Princeton, NJ), 307–317.
- Srivastava V, Neel J, MacKenzie A, Hicks J, DaSilva LA, Reed JH, Gilles R (2005) Using game theory to analyze wireless ad hoc networks. *IEEE Comm. Surveys and Tutorials* 7(4):46–56.
- Vetta A (2002) Nash equilibria in competitive societies with applications to facility location, traffic routing, and auctions. *Sympos. Foundations Comput. Sci.* (IEEE, Piscataway, NJ), 416–425.
- Villegas EG, Ferré RV, Josep Paradells J (2008) Frequency assignments in IEEE 802.11 WLANs with efficient spectrum sharing. *Wireless Comm. Mobile Comput.* 9(8):1125–1140.
- Wolpert D, Tumor K (1999) An overview of collective intelligence. Bradshaw JM, ed. *Handbook of Agent Technology* (AAAI Press/MIT Press, Cambridge, MA).
- Young HP (1994) *Equity* (Princeton University Press, Princeton, NJ).
- Zou Y, Chakrabarty K (2004) Uncertainty-award and coverage-oriented deployment for sensor networks. *J. Parallel and Distributing Comput.* 64(7):788–798.

Jason R. Marden is an assistant professor in the Department of Electrical, Computer, and Energy Engineering at the University of Colorado. In 2012, he received the Donald P. Eckman Award and an AFOSR Young Investigator Award. His research interests focus on game theoretic methods for feedback control of distributed multiagent systems.

Adam Wierman is a professor in the Department of Computing and Mathematical Sciences at the California Institute of Technology, where he is a member of the Rigorous Systems Research Group. His research interests center on resource allocation and scheduling decisions in computer systems and services. He received the 2011 ACM SIGMETRICS Rising Star award.