

You Get What You Share: Incentives for a Sharing Economy

Sreenivas Gollapudi
Google Research

Kostas Kollias
Google Research

Debmalya Panigrahi
Duke University

Abstract

In recent years, a range of online applications have facilitated resource sharing among users, resulting in a significant increase in resource utilization. In all such applications, sharing one's resources or skills with other agents increases social welfare. In general, each agent will look for other agents whose available resources complement hers, thereby forming natural *sharing groups*. In this paper, we study settings where a large population *self-organizes* into sharing groups. In many cases, centralized optimization approaches for creating an optimal partition of the user population are infeasible because either the central authority does not have the necessary information to compute an optimal partition, or it does not have the power to enforce a partition. Instead, the central authority puts in place an incentive structure in the form of a utility sharing method, before letting the participants form the sharing groups by themselves. We first analyze a simple *equal-sharing* method, which is the one most typically encountered in practice and show that it can lead to highly inefficient equilibria. We then propose a *Shapley-sharing* method and show that it significantly improves overall social welfare.

Introduction

The *sharing economy* (Belk 2014) is a term used to describe the modern trend of using online platforms and applications to increase the value of certain *resources* (such as goods, services, or skills) by enabling the sharing and reuse of these resources. Carpooling and ride sharing applications (e.g., Uber, Lyft, RideWith) allow people to better utilize commuting resources (cars, fuel, etc.), lodging and room sharing applications (e.g., AirBnB) allow users to share their space when not using it, online learning platforms (e.g., Coursera, Udacity, edX) allow students to share knowledge and collectively learn, and workforce applications (e.g., Upwork, Amazon Mechanical Turk) allow workers to collaboratively complete requested projects by contributing complementary skills. In all such applications, there are natural limits on the extent to which the resources can be shared. In carpooling, each car can fit a small number of people; in lodging, each room or house has a fixed capacity; in crowdsourced projects and online courses, there is a prescribed size for working groups. In effect, this means that the agents need to be organized in groups of limited size so as to effectively share

their individual resources. The platform allows the agents to list and search for available resources, which enables them to identify partners and form sharing groups. More specifically, in workforce and educational applications, there are groups that work to complete a task or a project, while in ride sharing and room sharing applications, the notion of a group appears when agents get together to use a provided resource, e.g., a ride or a house.

A natural goal is to partition the users into sharing groups that maximize the overall utility or social welfare of the system. One may model this as an optimization problem, where a centralized authority computes the optimal partition of users. However, in typical environments, the information necessary for computing this partition is hard to obtain due to privacy issues or due to the sheer size of the population. Moreover, the central authority has only limited control over the participants, and often cannot enforce the partition. Hence, in this paper, we study the following question: *Consider a population of individuals, who can form sharing groups and share their individual resources with group members. What is the right way to distribute the welfare produced by each sharing group among its members so that the population self-organizes in a socially optimal manner?* The semantics of distributing welfare are different for each application. For instance, in online learning, utility is given in the form of course credits that depend on how students contributed to group projects. In the ride-sharing application, using a vehicle has some value that is the same for all agents and monetary transfers reward those providing resources and extract payments from those utilizing resources to adjust utilities accordingly.

We model instances of this problem as games that we call *resource sharing games*. The most natural way to split the produced welfare among the group's members is to divide it equally. In fact, this is the method that naturally comes into play when there are no monetary or credit transfers in our applications of interest, i.e., everyone simply gets access to the resources available in the group. However, we show that this method can perform quite sub-optimally with respect to overall social welfare. We then propose an incentive structure derived from the well-known Shapley value in the economics literature (Shapley 1953), and show that the quality of the resulting outcomes is significantly improved.

In our model, each participant owns a subset of all possi-

ble resource types and, when a group is formed, everyone in the group has access to all resource types that at least one of them owns. An important observation is that the resulting social welfare function is not *submodular*, but *supermodular*. (This is in contrast to an existing body of work that studies utility games for submodular welfare functions, e.g., (Vetta 2002; Gairing 2009; Marden and Wierman 2013).) Informally, this means that a group produces more utility than the sum of its individuals. While this is often the case in practice, it can be theoretically problematic because the only stable solution is typically one where all the agents join the same group. However, in real-world applications (we give some examples below), it is impractical to have very large groups. So, we will adopt a model where groups are subject to a cardinality constraint, and also show that it is possible to enforce such a constraint through incentives alone.

Applications

First, we describe two concrete real world examples of our resource sharing games:

- *Massive Open Online Courses* – Suppose we are interested in partitioning the set of students into groups to complete an assignment comprising multiple problems. From a social perspective, we would prefer a partition of students such that every group in the partition completes as many problems as possible. Each student on the other hand, wishes to maximize the course credit she will receive. The goal of the central authority (teacher) is to provide the right *incentives* (in the form of a grading scheme) to the students so that effective groups grow organically and in a decentralized fashion.
- *Carpooling* – In a carpooling system, agents form groups and repeatedly share rides to and from work during rush hour. In this setting the resources of the game correspond to car rides from location to location and welfare is generated for each agent who takes a ride. From a social perspective, we wish that the commuting resources are put to good use. On the other hand, agents wish to optimize their individual rewards. Providing the appropriate incentives and inducing the organic formation of efficient sharing groups is important in achieving the social objective.

Several other applications such as crowdsourcing and labor markets similarly fit our proposed model.

Our Contributions

We consider a setting where a population of agents seeks to share resources in groups of size at most k . The value of k , the desired group size, is application specific (e.g., in (Davis 2009), values of 4, 5, and 6 are prescribed as appropriate for educational applications, while similar such constraints are natural in carpooling and other applications). Each group generates social welfare, which is then distributed as utility to the agents who generated it. We study this setting as a game, where the agents self-organize into groups. Each agent wishes to optimize her own utility share. We assume the role of the central authority, who cannot enforce the final sharing groups, but who can establish an incentive structure

(in the form of a utility-sharing method) that will lead the agents to efficient outcomes.

We begin by studying the most natural and widely used utility sharing method, which dictates that each agent has utility equal to the number of resources available in her sharing group. This is equivalent to equally splitting the social welfare produced by a group among its members and, hence, we call it *equal-sharing* (EQUAL). For instance, in online course projects, EQUAL would mean all students in a group will get the same course credit, i.e., the problems that the group has solved. In the carpooling application, EQUAL means all agents get to take the ride and split any costs (e.g., fuel) equally. We show that, for any reasonable equilibrium concept, there is an equilibrium partition of agents into sharing groups that is k times worse than the optimal. We also explain that this inefficiency with respect to the optimal is the worst possible for any partition, independent of whether it is an equilibrium or not. In light of this observation, we seek an alternative utility-sharing method, which results in better equilibria. The method we propose is derived from the Shapley value (Shapley 1953) and, hence, we call it *Shapley-sharing* (SHAPLEY). SHAPLEY aims to improve on the main shortcoming of EQUAL, which is that it rewards agents based on their contribution to improving the social welfare by sharing their resources. Below we present our results for the various equilibrium concepts, which show that SHAPLEY leads to better equilibria than EQUAL.

Nash equilibrium. The Nash equilibrium (NE) is the most established equilibrium concept in the literature. An NE in our setting is a partition such that no agent can improve her utility by unilaterally deviating to a different group. It is clear that the set of equilibria of a resource sharing game depends on the utility-sharing method being used. The *price of anarchy* (POA) (Koutsoupias and Papadimitriou 2009) is used as a metric that quantifies the performance of a utility-sharing method by calculating the worst-case ratio of the social welfare in the optimal partition to the social welfare in an NE. As we mentioned earlier, the POA of EQUAL is k . Unfortunately, the POA remains $\Theta(k)$ for SHAPLEY as well. In fact, we show that this will be the case for *any* natural utility-sharing method. However, we observe that the examples that yield this universal POA lower bound are unnatural as they assume agents will form groups of very small sizes, even when they are allowed to form larger sharing groups with more resources available. To get past these unnatural pathological cases, we consider the following two equilibrium concepts.

Strong Nash equilibrium. The strong Nash equilibrium (SNE) is an NE that is robust to deviations by coalitions of agents (Aumann 1959). This implies that a partition is an SNE when there is no subset of agents who can coordinate and deviate to other groups (or start groups on their own), thereby improving their individual utilities. The *strong price of anarchy* (SPOA) is the direct analog of the POA for the SNE. We prove that SHAPLEY results in a significant improvement of the SPOA, which is brought down to 2 (from k for EQUAL). Unfortunately, this equilibrium concept has the well-known drawback that it is not guaranteed to exist in many settings. Indeed, we give examples of resource shar-

ing games in our setting that have no SNE under SHAPLEY.

Toward the goal of characterizing equilibria with more desirable outcomes in the context of unilateral deviations, we focus on an equilibrium concept that we define here, and which we argue is the one that best predicts the behavior of selfish agents in our settings. We call this new concept a *balanced Nash equilibrium*.

Balanced Nash equilibrium. We define a *balanced Nash equilibrium* (BNE) as an NE such that all groups, except possibly one, have k agents, thereby ensuring balance in terms of their size. In many applications (such as in carpooling and online courses), this is an end in itself. In other settings, minimizing the number of groups subject to cardinality constraints is a desirable outcome. Note that merging small groups can only increase social welfare as more resources are added to the group and more agents benefit from them. We will also show that the BNE has the following properties:

- Ease of enforcement – It is easy for the central authority to slightly modify the SHAPLEY method so that every NE of the resulting game is a BNE of the original game, i.e., a BNE can be easily enforced.
- Existence – A BNE is always guaranteed to exist.
- Efficiency – The loss of efficiency in a BNE (we call this the *balanced price of anarchy* (BPOA) by analogy) is at most 3 in the SHAPLEY scheme (as against the BPOA of k of EQUAL).

Finally, we perform simulations which show that the SHAPLEY method improves on EQUAL in realistic settings, with respect to three metrics that measure our desiderata for forming sharing groups: *performance* (measured by the social welfare of the partition), *participation* (measured by the number of sharing groups of size k), and *fairness* (measured by the quality of the weakest group).

Preliminaries

In this section, we describe our model and introduce the notation. There is a set of *agents* N and a set of *resource types* P . Let $P_i \subseteq P$ denote the set of resource types that agent $i \in N$ possesses. Agents self-organize into groups so as to share their available resources. We denote the set of groups by \mathcal{G} . The social welfare of a group $G \in \mathcal{G}$ is:

$$U(G) = \begin{cases} |\bigcup_{i \in G} P_i| \cdot |G|, & \text{if } |G| \leq k \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where k is the *cardinality parameter* that represents the maximum group size. The expression encodes the fact that every agent in a group benefits from the presence of each resource available in the group. Note that this definition assumes that all resources have unit value. This assumption is without loss of generality since resources with non-uniform utility can be modeled by multiple resources with equal utility. The overall objective is to form sharing groups that maximize social welfare

$$U(\mathcal{G}) = \sum_{G \in \mathcal{G}} U(G).$$

Each agent i in group G , on the other hand, receives an individual utility $u_i(G)$. Utility-sharing is budget-balanced, i.e.,

$$\sum_{i \in G} u_i(G) = U(G).$$

We now define the equilibrium concepts that we will study.

Definition 1. A partition \mathcal{G} is a Nash equilibrium (NE) when for every agent i in group G , it is the case that $u_i(G) \geq u_i(G' \cup \{i\})$ for every other group $G' \in \mathcal{G}$. In other words, no agent can improve her utility by switching to any other sharing group unilaterally.

As mentioned earlier, we also use the coalition-resistant notion of strong Nash equilibrium.

Definition 2. A partition \mathcal{G} is a strong Nash equilibrium (SNE) if no coalition of agents $S \subseteq N$ can deviate (move to another group, including forming a new group) to form a new partition \mathcal{G}' such that $u_i(G') > u_i(G)$ for every agent $i \in S$, where G' and G are the respective groups of agent i in \mathcal{G}' and \mathcal{G} . Intuitively, an SNE is a partition such that no coalition of agents can coordinate and move between groups in a way that benefits every member of the coalition.

Finally we introduce the notion of a *balanced Nash equilibrium* (BNE).

Definition 3. A partition is a balanced Nash equilibrium (BNE) if it is an NE such that all groups, except possibly one, have k agents.

We will show existence of this equilibrium and demonstrate that it can be easily enforced by the central authority by modifying the utility-sharing method. We also note that BNE outcomes are desirable since they enforce groups of large size, where sharing is maximized and social welfare is increased.

Having introduced our three equilibrium concepts, we are ready to define the metrics that we use to measure the performance of a utility-sharing method.

Definition 4. The price of anarchy (POA), the strong price of anarchy (SPOA), and the balanced price of anarchy (BPOA) are the worst case ratios of the optimal social welfare to the social welfare of an NE, SNE, and BNE respectively, i.e.,

$$\begin{aligned} \text{POA} &= \frac{\max_{\mathcal{G}} U(\mathcal{G})}{\min_{\mathcal{G} \text{ is an NE}} U(\mathcal{G})}, \\ \text{SPOA} &= \frac{\max_{\mathcal{G}} U(\mathcal{G})}{\min_{\mathcal{G} \text{ is an SNE}} U(\mathcal{G})}, \\ \text{BPOA} &= \frac{\max_{\mathcal{G}} U(\mathcal{G})}{\min_{\mathcal{G} \text{ is a BNE}} U(\mathcal{G})}. \end{aligned}$$

Equal Sharing

The most natural way of sharing utility from resources in a group is to have every agent in the group receive utility equal to the value of the resources. Note that this is equivalent to equally splitting the welfare generated by the group among its participants and, hence, we call this utility-sharing method *equal-sharing* (EQUAL).

Theorem 5. *The POA, SPOA, and BPOA are all k under EQUAL.*

Proof. Consider the following setting. There is only one resource type. There are k agents who own the resource (we call them *strong* agents) and $k \cdot (k - 1)$ agents who do not (we call them *weak* agents). The optimal solution \mathcal{G}^* places every strong agent in a different group and adds $k - 1$ weak agents to each group. Then all agents have access to the unique resource type in the game and the social welfare is $U(\mathcal{G}^*) = k^2$. Now consider the partition \mathcal{G} where all strong agents are in the same group and all other agents are arbitrarily partitioned into groups of size k . We will show that this partition is an SNE, an NE, and a BNE. We first argue that it is an SNE: Note that all strong agents receive the maximum possible utility of 1, and hence they have no incentive to deviate from this partition in any way. Furthermore, the weak agents have no way to improve their utility above 0, since if any subset of them were to join the strong group, its cardinality would increase beyond k leading to total utility of 0. This proves that \mathcal{G} is an SNE. It follows that it is an NE as well, since every SNE is also an NE. Observing that the number of groups is equal to the ratio of the number of agents and the cardinality parameter, i.e., k , proves that \mathcal{G} is also a BNE. Note that $U(\mathcal{G}) = k$. Then we get that the POA, the SPOA, and the BPOA are all at least k .

It is not hard to see that this is actually the worst possible ratio for any possible partition (not only for equilibria). For each resource an agent owns, the least social welfare earned by that resource is 1, in cases when only that agent receives utility for it. The maximum possible is k , in cases when she shares it with another $k - 1$ agents in her group. \square

We note that this strong lower bound on the SPOA suggests that such inefficient equilibria can appear in many settings of interest, with weaker notions of stability. This includes situations where two agents can swap their group memberships or where a small number of agents can coordinate and deviate to a new group.

Shapley Sharing

In light of the poor performance of EQUAL in terms of the three inefficiency metrics, we seek a utility-sharing method that addresses the main shortcoming of EQUAL, which is that agents who significantly increase social welfare by sharing with their peers are not rewarded for it. The Shapley value (Shapley 1953) from the economics literature provides a method for splitting social welfare among those generating it. We propose applying the Shapley value in our setting, which yields a natural way of partially transferring utility from the agents who do not have a resource to the agent sharing it with them. In this section, we analyze the performance of *Shapley-sharing* (SHAPLEY) and show that it results in a significant improvement in the quality of equilibria.

The main idea is that when an agent is allowed access to a resource type that she does not possess, she shares her 1 unit of utility with agents who shared the resource with her.

Definition 6 (SHAPLEY). *SHAPLEY is based on the following simple utility granting rules:*

- *If agent i owns a resource of type p , then she gets 1 unit of utility for p .*
- *If agent i does not own a resource of type p and the group has n agents who have that resource and share it with her, then she gets $\frac{n}{n+1}$ units of utility for that resource. The remaining $\frac{1}{n+1}$ units are shared uniformly among those who shared resource p with her (so each of them gets $\frac{1}{n(n+1)}$ extra utility).*
- *If the size of the group is larger than k , each agent is penalized for a unit of utility per resource type.*

SHAPLEY is connected to the the standard definition of the Shapley value. The standard definition prescribes that each agent is rewarded with the expected increase she causes to the group welfare function $U(\cdot)$ over a uniformly random order of arrival of the agents in the group. In our setting the exact Shapley value we get when plugging in $U(\cdot)$ as the welfare functions is hard to work with and get a closed form expression. Our SHAPLEY instead is equivalent to the Shapley value under the following modification: Agents get utility equal to the expected increase they cause to the welfare function after removing the capacity constraint (i.e. the Shapley value with infinite capacity), but they also share a penalty equal to their cardinality if the size of the group is larger than k . This penalty is split evenly across agents since they all contribute equally to bringing the group size to larger than k .

SHAPLEY is an instantiation of the Shapley-value for the modified setting described above and, hence, guarantees the existence of an NE in every resource sharing game (Kollias and Roughgarden 2011). Intuitively, SHAPLEY is engineered to address the drawbacks that caused EQUAL to be very inefficient. It incentivizes strong agents to move to groups where they can offer more help and also incentivizes them to spread over multiple groups (since the bonus of an agent who shares is inversely proportional to the number of owners of the same resource type in the group). At the same time, it breaks symmetry in utility earned by the agents, allowing them to break away from very inefficient partitions that have groups of size k .

Nash Equilibria

By equivalence to Shapley value, we get that an NE is guaranteed to exist in our game for SHAPLEY utility sharing (Kollias and Roughgarden 2011). Indeed, we show in our full version (Gollapudi, Kollias, and Panigrahi 2018) that better-response dynamics always converges to an NE.

However, it turns out that any utility-sharing method, including SHAPLEY, that has an agent split her utility with those who shared resources with her has a POA of $\Theta(k)$. We call such methods *share-rewarding*. Due to space limitations, we omit the proof, which appears in our full version (Gollapudi, Kollias, and Panigrahi 2018).

Theorem 7. *The POA of any share-rewarding method is $\Theta(k)$.*

SHAPLEY is clearly a share-rewarding method, and hence it also has a POA of $\Theta(k)$.

Corollary 8. *The POA of SHAPLEY is $\Theta(k)$.*

Strong Nash Equilibria

In this section, we prove that the SPOA drops from k to 2 when we use SHAPLEY instead of EQUAL.

Theorem 9. *The SPOA of SHAPLEY is 2.*

Proof. We prove an upper bound of 2 on the SPOA of SHAPLEY here. In our full version (Gollapudi, Kollias, and Panigrahi 2018), we show that this analysis is tight, i.e., there is an example of a game where the SPOA of SHAPLEY is exactly 2.

Let \mathcal{G}^* be an optimal partition of agents into sharing groups, such that every group has at most k agents. Trivially one such optimal partition exists by the fact that any group with more than k agents produces 0 welfare by definition and, hence, breaking it into smaller groups can only improve the social welfare. Also, let \mathcal{G} be an SNE. Consider any group G^* of the optimal partition. We will show that the total utility of the agents from G^* in partition \mathcal{G} , for a given resource type p , is at least half the social welfare derived from this resource type in G^* .

Partition \mathcal{G} is an SNE, which implies there is at least one agent i_1 from group G^* , whose utility in her group in partition \mathcal{G} is at least her utility in group G^* . Otherwise, these agents would abandon their groups in partition \mathcal{G} and form G^* . If we remove agent i_1 from G^* , we get group $G^* \setminus \{i_1\}$. As before, there is at least one agent $i_2 \in G^* \setminus \{i_1\}$ whose utility in partition \mathcal{G} is at least her utility in group $G^* \setminus \{i_1\}$ because otherwise the agents in $G^* \setminus \{i_1\}$ would deviate to form $G^* \setminus \{i_1\}$. Continue in the same fashion until we have a complete ordering $i_1, i_2, \dots, i_{|G^*|}$ of G^* , such that the utility of i_l in her group in \mathcal{G} is at least her utility in group $\{i_l, i_{l+1}, \dots, i_{|G^*|}\}$.

Let u be the sum over all $l = 1, 2, \dots, |G^*|$, of the utility of agent i_l , in group $\{i_l, i_{l+1}, \dots, i_{|G^*|}\}$. As argued, this sum lower bounds the total utility of these agents in \mathcal{G} . Let u_p denote the contribution of a resource type p to the value of u . Now focus on an agent i_l and examine the utility she gets from resource type p . There are three cases:

- If she owns p and there is also some $i_{l'}$ who owns p , with $l' > l$, then the contribution of i_l to u_p is at least 1, since she gets 1 unit of utility for owning p and we have assumed $|G^*| \leq k$.
- If she does not own p , but there is some $i_{l'}$ who owns p , with $l' > l$, then the contribution of i_l to u_p is at least $\frac{1}{2}$, since she gets to use it once $i_{l'}$ shares it, in which case she gets at least $\frac{1}{2}$ utility.
- If i_l owns p , and she is the last one in our ordering of G^* who does, then her contribution to u_p is at least 1 plus $\frac{1}{2}$ for every agent after her in the ordering, since she gets 1 unit of utility for owning p and $\frac{1}{2}$ for every agent she shares it with in $\{i_l, \dots, i_{|G^*|}\}$.

From the above, it follows that we get at least $\frac{1}{2}$ in u_p for every agent in G^* . The optimal partition extracts 1 unit of social welfare from each agent in G^* for p . Summing over all resource types and over all groups of \mathcal{G}^* , and using the fact that the SHAPLEY method distributes precisely the social welfare as utility, we get the SPOA upper bound of 2. \square

While this bound shows that SHAPLEY outperforms EQUAL for the SPOA metric, an SNE may not exist in all games – we give such an example in our full version (Gollapudi, Kollias, and Panigrahi 2018).

Fact 10. *There exist games that do not have an SNE under SHAPLEY sharing.*

Balanced Nash Equilibria

We have previously seen that SHAPLEY outperforms EQUAL for an SNE if it exists, but an SNE may not exist in some games. On the other hand, an NE exists in all games, but all natural utility sharing methods (share-rewarding methods) are highly inefficient in this case. So, the question arises: is there an equilibrium concept that can be guaranteed to exist and has efficiency comparable to SNE? We show that a BNE achieves these properties.

Recall that a BNE is an NE with $\left\lceil \frac{|N|}{k} \right\rceil$ groups, such that all groups, except possibly one, have exactly k agents each, where N is the set of agents. We note that the central authority can easily enforce such a partition, by prescribing $\left\lceil \frac{|N|}{k} \right\rceil$ sharing groups that the agents should join and by assigning negative utility to the agents not participating in the prescribed groups. Moreover, the authority needs to enforce a cardinality constraint of k for the first $\left\lfloor \frac{|N|}{k} \right\rfloor$ groups and a cardinality constraint of $|N| - k \left\lfloor \frac{|N|}{k} \right\rfloor$ for the last group. We show that with SHAPLEY, we can bring the BPOA from k (in EQUAL) down to at most 3. Our BPOA bound can be read both as a BPOA bound for the original game and as a POA bound for the modified game that fixes the number of sharing groups by giving negative utility to non-prescribed groups.

A BNE in the original game is guaranteed to exist, as exhibited by the following theorem, which we prove in our full version (Gollapudi, Kollias, and Panigrahi 2018).

Theorem 11. *SHAPLEY always induces games such that a BNE exists.*

We now prove our upper bound on the BPOA. The intuition behind the fact that the BPOA improves significantly, in contrast to the POA, is that almost all groups have a large number of agents and there is an incentive for strong agents to move to a group that has few resource types available and receive credit for sharing resources.

Theorem 12. *The BPOA of SHAPLEY is at most 3.*

Proof. For simplicity of exposition, in this proof we assume that the number of agents is a multiple of k . Our arguments easily extend to the general case.

Let \mathcal{G} be a BNE partition with $g = \frac{|N|}{k}$ sharing groups. For the purposes of this proof we define a new notion that we call the *uniform fractional deviation* of an agent from \mathcal{G} . In such a deviation, the agent splits herself into g fractions of size $\frac{1}{g}$ each and allocates one such fraction to each sharing group. When an agent picks this strategy, her utility is her average utility across all groups, i.e., $\sum_{G \in \mathcal{G}} \frac{u_i(G \cup \{i\})}{g}$, $u_i(G)$ be-




Figure 1: SHAPLEY yields $\sim 10\%$ improvement of the social welfare across agent, consistently fills all groups, and significantly improves the worst agent.

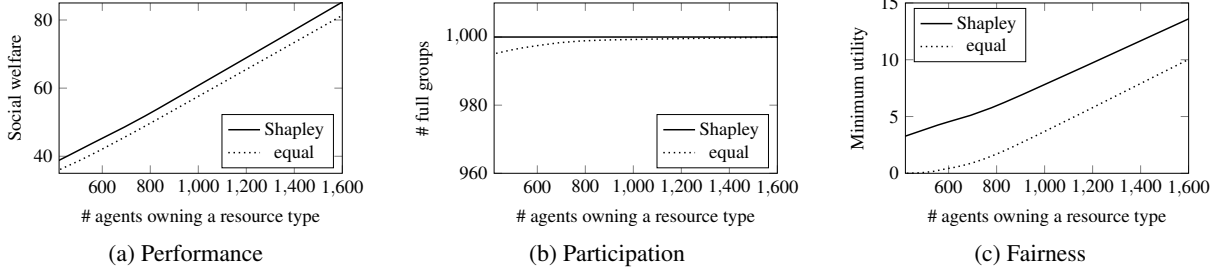


Figure 2: SHAPLEY yields $\sim 10\%$ improvement of the social welfare across resource type frequency levels, consistently fills all sharing groups, and significantly improves the minimum utility.

ing the utility of i in group G . When all agents play this strategy, we get the *uniform fractional partition* \mathcal{G}^* . The welfare of the uniform fractional partition is given as follows: Suppose n_p agents own a resource of type p . Then there is a cumulative of $\frac{n_p}{g}$ agents who own this resource type in each group, since each group has a total cumulative fraction of k agents. The welfare extracted from resource type p in \mathcal{G}^* is $\min\left\{1, \frac{n_p}{g}\right\}kg = \min\{gk, n_pk\}$, which is the best possible.

We now prove that the total welfare in the uniform fractional partition \mathcal{G}^* is at most 3 times the welfare of the BNE \mathcal{G} . We do so by applying the (λ, μ) -smoothness framework (Roughgarden 2009) which prescribes that, if we find numbers λ and μ such that

$$\sum_{i \in N} \sum_{G \in \mathcal{G}} \frac{u_i(G \cup \{i\})}{g} \geq \lambda U(\mathcal{G}^*) - \mu U(\mathcal{G}), \quad (2)$$

then the BPOA is at most $\frac{\mu+1}{\lambda}$. We may see this as follows: By the fact that \mathcal{G} is a BNE, we get that each agent i does not benefit by switching to any sharing group other than her current one. Hence, i also does not benefit by performing the uniform fractional deviation, since the resulting utility is the average of all possible integral deviations. Then, it follows that, $U(\mathcal{G}) \geq \lambda U(\mathcal{G}^*) - \mu U(\mathcal{G})$, which, by rearranging, gives a $\frac{\mu+1}{\lambda}$ upper bound on the BPOA.

Clearly, it suffices to find λ and μ such that inequality (2) is true per resource type, i.e., if the agent's deviation utility on the left hand side only takes a given resource type p into consideration and the social welfare on the right hand side, again only takes resource type p into consideration. Then,

summing over all resource types would give us (2). Given the above, we will show that $\lambda = \mu = \frac{1}{2}$ satisfy the inequality on a per resource type basis, thus proving an upper bound of $\frac{\mu+1}{\lambda} = 3$ on the BPOA.

Suppose \mathcal{G}^* yields social welfare $U_p(\mathcal{G}^*)$ as far as resource type p is concerned and that \mathcal{G} yields social welfare $U_p(\mathcal{G}) = \gamma U_p(\mathcal{G}^*)$. We assume $\gamma \in [0, 1]$, since, if this is not the case, we can discard this resource type from our comparison between \mathcal{G} and \mathcal{G}^* , without loss of generality. There are $|N| - \gamma U_p(\mathcal{G}^*)$ agents that do not own resource type p in \mathcal{G} . Hence, when an agent who owns the corresponding resource type deviates to her uniform fractional strategy, she will get utility of at least $\frac{k}{|N|} \frac{|N| - \gamma U_p(\mathcal{G}^*)}{2}$, since $\frac{k}{|N|}$ is the size of the fraction she assigns to each sharing group and $\frac{1}{2}$ is the reward she gets for helping the $|N| - \gamma U_p(\mathcal{G}^*)$ integral agents who do not own resource type p . There are at least $\frac{U_p(\mathcal{G}^*)}{k}$ agents who own resource type p , given that \mathcal{G}^* has social welfare $U_p(\mathcal{G}^*)$ for p . Then we get that the part of the left-hand side of (2) that corresponds to p is

$$\frac{k}{|N|} \frac{|N| - \gamma U_p(\mathcal{G}^*)}{2} \frac{U_p(\mathcal{G}^*)}{k}.$$

If we substitute $|N|$ with its smaller or equal number $U_p(\mathcal{G}^*)$ we get that the above is at most $\frac{U_p(\mathcal{G}^*)}{2} (1 - \gamma)$, for which the inequality $\frac{U_p(\mathcal{G}^*)}{2} (1 - \gamma) \geq \lambda U_p(\mathcal{G}^*) - \mu \gamma U_p(\mathcal{G}^*)$ always holds with $\lambda = \mu = \frac{1}{2}$. This completes the proof. \square

Simulations

In this section, we experimentally evaluate the performance of SHAPLEY vs EQUAL. In our setup, we begin with $|N| = 5,000$ agents and a hypothetical set of $|P| = 20$ resource types that all agents can utilize in groups of size at most $k = 5$. In each simulation, we begin with each agent in a sharing group by herself. Then, in each round we let the agents deviate (one by one) to a sharing group that they prefer until no such deviations are possible, i.e., until an NE is reached.

We use three metrics to evaluate the quality of the partition reached by the agents. The first one is the natural *performance* metric, which measures the social welfare of the resulting partition. The second one is a *participation* metric, which counts the number of sharing groups for which the cardinality constraint is tight. As we will see, SHAPLEY almost always yields partitions such that all groups are full, which also validates our study of the BNE as an important equilibrium concept in this setting. Finally, we focus on a *fairness* metric, which measures the minimum social welfare extracted by an agent, i.e., the number of resources available in the weakest group.

In the following sets of experiments we use skewed distributions that follow power laws. This is to model the fact that in most applications there is a small fraction of agents with many resources and most agents have few resources (e.g., in ride sharing settings most agents contribute a single car, with few having more than that and a tiny fraction operating a fleet of cars). Similarly, there are a few types of resources that are very popular. We present the results for two datasets below, and give additional simulation results in our full version (Gollapudi, Kollias, and Panigrahi 2018). All the simulation results exhibit similar trends, where SHAPLEY improves significantly over EQUAL with respect to the parameters described above.

Skewed Agent Resource Ownage. For our first simulation, we assume that the 20 resource type frequencies are symmetric (i.e., we expect approximately the same number of agents will own each one). On the contrary, we assume that the agent resource ownage, defined as the number of resource types the agent owns, follows a power law distribution. The plots in Figures 1a, 1b, and 1c show how our three metrics vary as the average ownage level changes (i.e., as the degree of the power law changes).

Performance-wise, we see in Figure 1a that SHAPLEY improves the total utility by approximately 10% for all agent ownage levels. We can also see (from Figure 1b) that agents consistently participate in sharing groups when SHAPLEY is applied. In the proof of Theorem 12, we showed that when the number of groups is large, the BPOA of SHAPLEY is upper bounded by 3. On the other hand, in Theorem 5, we showed that the POA of EQUAL is $k = 5$. The improvements from SHAPLEY indeed verify our theoretical analysis in this setting.

Figure 1b shows that EQUAL can have unfilled groups as the average agent ownage level drops, since there are more agents who cannot access any resource type. Since the other agents are clustered in groups, they are left out and have no incentive to join each other. For the same reason, EQUAL performs poorly with respect to minimum utility, as we see

in Figure 1c. In contrast, SHAPLEY performs better on this metric by rewarding agents for sharing resources.

Skewed Resource Type Frequency. For our second simulation, we assume that the agents are symmetric with respect to their resource ownage expressed as a fraction of resources owned. Further, we assume that the frequency of each resource type, defined as the number of agents who own the corresponding resource type, is drawn from a power law distribution. The plots in Figures 2a, 2b, and 2c show how our three metrics vary as the average resource type frequency changes (i.e., as the degree of the power law changes).

Similar to the previous simulation, we observe in Figure 2b that SHAPLEY results in balanced partitions. From the plot of Figure 2a, we conclude that SHAPLEY offers a significant improvement in the social welfare for this setting as well. These two facts are further validation of our theoretical analyses.

In summary, SHAPLEY outperforms EQUAL on all three metrics. This is expected, given that now the agents are considered symmetric and there are only a few agents who do not own any resource types. This also results in the minimum utility to be above 0 and the number of filled sharing groups to increase.

Related Work

The price of anarchy was defined in (Koutsoupias and Papadimitriou 2009) (we point the reader to (Nisan et al. 2007) for the main results of the area). The study of the inefficiency of further equilibrium concepts naturally followed, including the strong Nash equilibrium (Aumann 1959; Andelman, Feldman, and Mansour 2009; Bachrach et al. 2013; Roughgarden and Schrijvers 2014).

Vetta (Vetta 2002) considered a broad class of utility games and gave an upper bound of 2 on the price of anarchy. The interpretation of the main assumption of (Vetta 2002) in our setting is that the social welfare function in each sharing group is submodular. While this is a reasonable assumption in many settings, this does not hold in the applications we study, where submodular utility functions do not properly model realistic scenarios and lead to agents forming singleton groups. Following up on (Vetta 2002), other more specific applications have been considered in (Marden and Wierman 2013; Marden and Roughgarden 2014) as also a setting with incomplete information (Bachrach, Syrgkanis, and Vojnovic 2013).

Coalitionally robust resource sharing games can be placed in the context of hedonic games, a term coined by Dreze and Greenberg (Dreze and Greenberg 1980). The earliest and most well-known result in the area is the one on stable marriages by Gale and Shapley (Gale and Shapley 1962). Further works include (Aumann and Dreze 1974; Apt et al. 2014; Banerjee, Konishi, and Sonmez 2001; Bogomolnaia and Jackson 2002; Chalkiadakis et al. 2009; Feldman, Lewin-Eytan, and Naor 2012; Greenberg and Weber 1993; Hajdukova 2006; Hart and Kurz 1983; Immorlica, Markakis, and Piliouras 2010; Yi 1997).

The Shapley value, upon which our SHAPLEY method is based, is defined in (Shapley 1953). The Shapley value

has been used in utility games (Marden and Wierman 2013; Marden and Roughgarden 2014) and in network routing games (Kollias and Roughgarden 2011). The strong price of anarchy of games induced by the Shapley value in network cost-sharing games is studied in (Roughgarden and Schrijvers 2014).

Conclusion

In this paper, we assumed the central authority's perspective in sharing economy settings where agents organically form sharing groups to maximize their individual utility. We showed that a Nash equilibrium (NE) is always guaranteed to exist, but all natural utility sharing methods are inefficient for the corresponding price of anarchy (POA) metric. In contrast, a Shapley utility sharing method (SHAPLEY) is highly efficient for the strong POA metric, and significantly outperforms the commonly used equal utility sharing (EQUAL) method. However, the corresponding equilibrium concept of a strong NE may not exist for all instances. We combined the benefits of these solution concepts by defining a new equilibrium concept called balanced NE, which always exists and is nearly as efficient as strong NE for Shapley sharing. Simulations confirm that SHAPLEY outperforms EQUAL not only in terms of social welfare but also for other desirable metrics such as participation and fairness. An interesting future direction is to study other welfare functions such as the ratio of social welfare to the sum of available resources, which measures the extent to which the agents are helped by their peers beyond what they already own.

References

- Andelman, N.; Feldman, M.; and Mansour, Y. 2009. Strong price of anarchy. *Games and Economic Behavior* 289–317.
- Apt, K. R.; Rahn, M. M.; Schafer, G.; and Simon, S. E. 2014. Coordination games on graphs. In *Conference on Web and Internet Economics 2014 (WINE 14)*.
- Aumann, R. J., and Dreze, J. H. 1974. Cooperative games with coalition structures. *International Journal of Game Theory* 3(4):217–237.
- Aumann, R. J. 1959. Acceptable points in general cooperative n-person games. *Contribution to the theory of game IV, Annals of Mathematical Study* 40:287–324.
- Bachrach, Y.; Syrgkanis, V.; Tardos, É.; and Vojnovic, M. 2013. Strong price of anarchy and coalitional dynamics. *CoRR* abs/1307.2537.
- Bachrach, Y.; Syrgkanis, V.; and Vojnovic, M. 2013. Incentives and efficiency in uncertain collaborative environments. In *Web and Internet Economics - 9th International Conference, WINE 2013, Cambridge, MA, USA, December 11-14, 2013, Proceedings*, 26–39.
- Banerjee, S.; Konishi, H.; and Sonmez, T. 2001. Core in a simple coalition formation game. *Social Choice and Welfare* 135–153.
- Belk, R. 2014. You are what you can access: Sharing and collaborative consumption online. *Journal of Business Research* 67(8):1595–1600.
- Bogomolnaia, A., and Jackson, M. O. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38(2):201–230.
- Chalkiadakis, G.; Elkind, E.; Polukarov, M.; and Jennings, N. R. 2009. The price of democracy in coalition formation. In *8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, 401–408.
- Davis, B. G. 2009. *Tools for Teaching*.
- Dreze, J. H., and Greenberg, J. 1980. Hedonic coalitions: Optimality and stability. *Econometrica* 48(4):987–1003.
- Feldman, M.; Lewin-Eytan, L.; and Naor, J. 2012. Hedonic clustering games. In *24th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '12, Pittsburgh, PA, USA, June 25-27, 2012*, 267–276.
- Gairing, M. 2009. Covering games: Approximation through non-cooperation. In *WINE*, 184–195.
- Gale, D., and Shapley, L. S. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69(1):9–15.
- Gollapudi, S.; Kollias, K.; and Panigrahi, D. 2018. You get what you share: Incentives for a sharing economy. <http://sreenivasgollapudi.com/pubs/AAAI2019.pdf>.
- Greenberg, J., and Weber, S. 1993. Stable coalition structures with a unidimensional set of alternatives. *J. Econ. Theory* 60:62–82.
- Hajdukova, J. 2006. Coalition formation games: a survey. *IGTR* 8(4):613–641.
- Hart, S., and Kurz, M. 1983. Endogenous formation of coalitions. *Econometrica* 51(4):1047–1064.
- Immorlica, N.; Markakis, E.; and Piliouras, G. 2010. Coalition formation and price of anarchy in cournot oligopolies. In *WINE*, 270–281.
- Kollias, K., and Roughgarden, T. 2011. Restoring pure equilibria to weighted congestion games. In *ICALP*.
- Koutsoupias, E., and Papadimitriou, C. H. 2009. Worst-case equilibria. *Computer Science Review* 3(2):65–69.
- Marden, J. R., and Roughgarden, T. 2014. Generalized efficiency bounds in distributed resource allocation. *IEEE Trans. Automat. Contr.* 59(3):571–584.
- Marden, J. R., and Wierman, A. 2013. Distributed welfare games. *Operations Research* 61(1):155–168.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic Game Theory*.
- Roughgarden, T., and Schrijvers, O. 2014. Network cost-sharing without anonymity. In *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*, 134–145.
- Roughgarden, T. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, 513–522.
- Shapley, L. S. 1953. *Additive and Non-Additive Set Functions*. Ph.D. Dissertation, Department of Mathematics, Princeton University.
- Vetta, A. 2002. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS*.
- Yi, S. 1997. Stable coalition structures with externalities. *Games and Economic Behavior* 20:201–237.