
Statistical Cost Sharing

Eric Balkanski
Harvard University
ericbalkanski@g.harvard.edu

Umar Syed
Google NYC
usyed@google.com

Sergei Vassilvitskii
Google NYC
sergeiv@google.com

Abstract

We study the cost sharing problem for cooperative games in situations where the cost function C is not available via oracle queries, but must instead be learned from samples drawn from a distribution, represented as tuples $(S, C(S))$, for different subsets S of players. We formalize this approach, which we call STATISTICAL COST SHARING, and consider the computation of the core and the Shapley value. Expanding on the work by Balcan et al. [2015], we give precise sample complexity bounds for computing cost shares that satisfy the core property with high probability for any function with a non-empty core. For the Shapley value, which has never been studied in this setting, we show that for submodular cost functions with bounded curvature κ it can be approximated from samples from the uniform distribution to a $\sqrt{1 - \kappa}$ factor, and that the bound is tight. We then define statistical analogues of the Shapley axioms, and derive a notion of statistical Shapley value and that these can be approximated arbitrarily well from samples from any distribution and for any function.

1 Introduction

The cost sharing problem asks for an equitable way to split the cost of a service among all of the participants. Formally, there is a cost function C defined over all subsets $S \subseteq N$ of a ground set of elements, or *players*, and the objective is to fairly divide the cost of the ground set $C(N)$ among the players. Unlike traditional learning problems, the goal here is not to predict the cost of the service, but rather learn which ways of dividing the cost among the players are equitable.

Cost sharing is central to cooperative game theory, and there is a rich literature developing the key concepts and principles to reason about this topic. Two popular cost sharing concepts are the *core* [Gillies, 1959], where no group of players has an incentive to deviate, and the *Shapley value* [Shapley, 1953], which is the unique vector of cost shares satisfying four natural axioms.

While both the core and the Shapley value are easy to define, computing them poses additional challenges. One obstacle is that the computation of the cost shares requires knowledge of costs in myriad different scenarios. For example, computing the exact Shapley value requires one to look at the marginal contribution of a player over *all possible subsets* of others. Recent work [Liben-Nowell et al., 2012] shows that one can find approximate Shapley values for a restricted subset of cost functions by looking at the costs for polynomially many specifically chosen subsets. In practice, however, another roadblock emerges: one cannot simply query for the cost of an arbitrary subset. Rather, the subsets are passively observed, and the costs of unobserved subsets are simply unknown. We share the opinion of Balcan et al. [2016] that the main difficulty with using cost sharing methods in concrete applications is the information needed to compute them.

Concretely, consider the following cost sharing applications.

Attributing Battery Consumption on Mobile Devices. A modern mobile phone or tablet is typically running a number of distinct apps concurrently. In addition to foreground processes, a lot of activity may be happening in the background: email clients may be fetching new mail, GPS may be active for geo-fencing applications, messaging apps are polling for new notifications, and so on. All of these activities consume power; the question is how much of the total battery consumption should be attributed to each app? This problem is non-trivial because the operating system induces cooperation between apps to save battery power. For example there is no need to activate the GPS sensor twice if two different apps request the current location almost simultaneously.

Understanding Black Box Learning Deep neural networks are prototypical examples of black box learning, and it is almost impossible to tease out the contribution of a particular feature to the final output. Particularly in situations where the features are binary, cooperative game theory gives a formal way to analyze and derive these contributions. While one can evaluate the objective function on any subset of features, deep networks are notorious for performing poorly on certain out of sample examples [Goodfellow et al., 2014, Szegedy et al., 2013], which may lead to misleading conclusions when using traditional cost sharing methods.

We model these cost sharing questions as follows. Let N be the set of possible players (apps or features), and for a subset $S \subseteq N$, let $C(S)$ denote the cost of S . This cost represents the total power consumed over a standard period of time, or the rewards obtained by the learner. We are given ordered pairs $(S_1, C(S_1)), (S_2, C(S_2)), \dots, (S_m, C(S_m))$, where each $S_i \subseteq N$ is drawn independently from some distribution \mathcal{D} . The problem of STATISTICAL COST SHARING asks to look for reasonable cost sharing strategies in this setting.

1.1 Our results

We build on the approach from Balcan et al. [2015], which studied STATISTICAL COST SHARING in the context of the core, and assume that only partial data about the cost function is observed. The authors showed that cost shares that are likely to respect the core property can be obtained for certain restricted classes of functions. Our main result is an algorithm that generalizes these results for *all* games where the core is non-empty and we derive sample complexity bounds showing exactly the number of samples required to compute cost shares (Theorems 1 and 2). While the main approach of Balcan et al. [2015] relied on first learning the cost function and then computing cost shares, we show how to proceed directly, computing cost shares without explicitly learning a good estimate of the cost function. This high level idea was independently discovered by Balcan et al. [2016]; our approach here greatly improves the sample complexity bounds, culminating in a result logarithmic in the number of players. We also show that approximately satisfying the core with probability one is impossible in general (Theorem 3).

We then focus on the Shapley value, which has never been studied in the STATISTICAL COST SHARING context. We obtain a tight $\sqrt{1 - \kappa}$ multiplicative approximation of the Shapley values for submodular functions with bounded curvature κ over the uniform distribution (Theorems 4 and 11), but show that they cannot be approximated by a bounded factor in general, even for the restricted class of coverage functions, which are learnable, over the uniform distribution (Theorem 5). We also introduce a new cost sharing method called *data-dependent Shapley value* which is the unique solution (Theorem 6) satisfying four natural axioms resembling the Shapley axioms (Definition 7), and which can be approximated arbitrarily well from samples for any bounded function and any distribution (Theorem 7).

1.2 Related work

There are two avenues of work which we build upon. The first is the notion of cost sharing in cooperative games, first introduced by Von Neumann and Morgenstern [1944]. We consider the Shapley value and the core, two popular solution concepts for cost-sharing in cooperative games. The Shapley value [Shapley, 1953] is studied in algorithmic mechanism design [Anshelevich et al., 2008, Balkanski and Singer, 2015, Feigenbaum et al., 2000, Moulin, 1999]. For applications of the Shapley value, see the surveys by Roth [1988] and Winter [2002]. A naive computation of the Shapley value of a cooperative game would take exponential time; recently, methods for efficiently approximating

the Shapley value have been suggested [Bachrach et al., 2010, Fatima et al., 2008, Liben-Nowell et al., 2012, Mann, 1960] for some restricted settings.

The core, introduced by Gillies [1959], is another well-studied solution concept for cooperative games. Bondareva [1963] and Shapley [1967] characterized when the core is non-empty. The core has been studied in the context of multiple combinatorial games, such as facility location Goemans and Skutella [2004] and maximum flow Deng et al. [1999]. In cases with no solutions in the core or when it is computationally hard to find one, the balance property has been relaxed to hold approximately [Devanur et al., 2005, Immorlica et al., 2008]. In applications where players submit bids, cross-monotone cost sharing, a concept stronger than the core that satisfies the group strategy proofness property, has attracted a lot of attention [Immorlica et al., 2008, Jain and Vazirani, 2002, Moulin and Shenker, 2001, Pál and Tardos, 2003]. We note that these applications are sufficiently different from the ones we are studying in this work.

The second is the recent work in econometrics and computational economics that aims to estimate critical concepts directly from a limited data set, and reason about the sample complexity of the computational problems. Specifically, in all of the above papers, the algorithm must be able to query or compute $C(S)$ for an arbitrary set $S \subseteq N$. In our work, we are instead given a collection of samples from some distribution; importantly the algorithm does not know $C(S)$ for sets S that were not sampled. This approach was first introduced by Balcan et al. [2015], who showed how to compute an approximate core for some families of games. Their main technique is to first learn the cost function C from samples and then to use the learned function to compute cost shares. The authors also showed that there exist games that are not PAC-learnable but that have an approximate core that can be computed. Independently, in recent follow up work, the authors showed how to extend their approach to compute a probably approximate core for all games with a non-empty core, and gave weak sample complexity bounds [Balcan et al., 2016]. We improve upon their bounds, showing that a logarithmic number of samples suffices when the spread of the cost function is bounded.

2 Preliminaries

A *cooperative game* is defined by an ordered pair (N, C) , where N is the ground set of *elements*, also called *players*, and $C : 2^N \rightarrow \mathbb{R}_{\geq 0}$ is the *cost function* mapping each *coalition* $S \subseteq N$ to its cost, $C(S)$. The ground set of size $n = |N|$ is called the *grand coalition* and we denote the elements by $N = \{1, \dots, n\} = [n]$. We assume that $C(\emptyset) = 0$, $C(S) \geq 0$ for all $S \subseteq N$, and that $\max_S C(S)$ is bounded by a polynomial in n , which are standard assumptions. We will slightly abuse notation and use $C(i)$ instead of $C(\{i\})$ for $i \in N$ when it is clear from the context.

We recall three specific classes of functions. *Submodular* functions exhibit the property of diminishing returns: $C_S(i) \geq C_T(i)$ for all $S \subseteq T \subseteq N$ and $i \in N$ where $C_S(i)$ is the marginal contribution of element i to set S , i.e., $C_S(i) = C(S \cup \{i\}) - C(S)$. *Coverage* functions are the canonical example of submodular functions. A function is coverage if it can be written as $C(S) = |\cup_{i \in S} T_i|$ where $T_i \subseteq U$ for some universe U . Finally, we also consider the simple class of additive functions, such that $C(S) = \sum_{i \in S} C(i)$.

A *cost allocation* is a vector $\psi \in \mathbb{R}^n$ where ψ_i is the *share* of element i . We call a cost allocation ψ *balanced* if $\sum_{i \in N} \psi_i = C(N)$. Given a cooperative game (N, C) the goal in the cost sharing literature is to find “desirable” balanced cost allocations. Most proposals take an axiomatic approach, defining a set of axioms that a cost allocation should satisfy. These lead to the concepts of Shapley value and the core, which we define next. A useful tool to describe and compute these cost sharing concepts is permutations. We denote by σ a uniformly random permutation of N and by $S_{\sigma < i}$ the players before $i \in N$ in permutation σ .

2.1 The core

The core is a balanced cost allocation where no player has an incentive to deviate from the grand coalition — for any subset of players the sum of their shares does not cover their collective cost.

Definition 1. A cost allocation ψ is in the core of function C if the following properties are satisfied:

- **Balance:** $\sum_{i \in N} \psi_i = C(N)$,
- **Core property:** for all $S \subseteq N$, $\sum_{i \in S} \psi_i \leq C(S)$.

The core is a natural cost sharing concept. For example, in the battery blame scenario it translates to the following assurance: No matter what other apps are running concurrently, an app is never blamed for more battery consumption than if it were running alone. Given that app developers are typically business competitors, and that a mobile device’s battery is a very scarce resource, such a guarantee can rather neatly avoid a great deal of finger-pointing. Unfortunately, for a given cost function C the core may not exist (we say the core is empty), or there may be multiple (or even infinitely many) cost allocations in the core. For submodular functions C , the core is guaranteed to exist and one allocation in the core can be computed in polynomial time. Specifically, for any permutation σ , the cost allocation ψ such that $\psi_i = C(S_{\sigma < i} \cup \{i\}) - C(S_{\sigma < i})$ is in the core.

2.2 The Shapley value

The Shapley value provides an alternative cost sharing method. For a game (N, C) we denote it by ϕ^C , dropping the superscript when it is clear from the context. While the Shapley value may not satisfy the core property, it satisfies the following four axioms:

- **Balance:** $\sum_{i \in N} \phi_i = C(N)$.
- **Symmetry:** For all $i, j \in N$, if $C(S \cup \{i\}) = C(S \cup \{j\})$ for all $S \subseteq N \setminus \{i, j\}$ then $\phi_i = \phi_j$.
- **Zero element:** For all $i \in N$, if $C(S \cup \{i\}) = C(S)$ for all $S \subseteq N$ then $\phi_i = 0$.
- **Additivity:** For two games (N, C_1) and (N, C_2) with the same players, but different cost functions C_1 and C_2 , let ϕ^1 and ϕ^2 be the respective cost allocations. Consider a new game $(N, C_1 + C_2)$, and let ϕ' be the cost allocation for this game. Then for all elements, $i \in N$, $\phi'_i = \phi_i^1 + \phi_i^2$.

Each of these axioms is natural: balance ensures that the cost of the grand coalition is distributed among all of the players. Symmetry states that two identical players should have equal shares. Zero element verifies that a player that adds zero cost to any coalition should have zero share. Finally, additivity just confirms that costs combine in a linear manner. It is surprising that the set of cost allocations that satisfies all four axioms is unique. Moreover, the Shapley value ϕ can be written as the following summation:

$$\phi_i = \mathbb{E}_{\sigma} [C(S_{\sigma < i} \cup \{i\}) - C(S_{\sigma < i})] = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n - |S| - 1)!}{n!} (C(S \cup \{i\}) - C(S)).$$

This expression is the expected marginal contribution $C(S \cup \{i\}) - C(S)$ of i over a set of players S who arrived before i in a random permutation of N . As the summation is over exponentially many terms, the Shapley value generally cannot be computed exactly in polynomial time. However, several sampling approaches have been suggested to approximate the Shapley value for specific classes of functions Bachrach et al. [2010], Fatima et al. [2008], Liben-Nowell et al. [2012], Mann [1960].

2.3 Statistical cost sharing

With the sole exception of Balcan et al. [2015], previous work in cost-sharing critically assumes that the algorithm is given oracle access to C , i.e., it can query, or determine, the cost $C(S)$ for any $S \subseteq N$. In this paper, we aim to (approximately) compute the Shapley value and other cost allocations from *samples*, without oracle access to C , and with a number of samples that is polynomial in n .

Definition 2. Consider a cooperative game with players N and cost function C . In the STATISTICAL COST SHARING problem we are given pairs $(S_1, C(S_1)), (S_2, C(S_2)), \dots, (S_m, C(S_m))$ where each S_i is drawn i.i.d. from a distribution \mathcal{D} over 2^N . The goal is to find a cost allocation $\psi \in \mathbb{R}^n$.

In what follows we will often refer to an individual $(S, C(S))$ pair as a *sample*. It is tempting to reduce STATISTICAL COST SHARING to classical cost sharing by simply collecting enough samples to use known algorithms. For example, Liben-Nowell et al. [2012] showed how to approximate the Shapley value with polynomially many queries $C(S)$. However, if the distribution \mathcal{D} is not aligned with these specific queries, which is the case even for the uniform distribution, emulating these

algorithms in our setting requires exponentially many samples. Balcan et al. [2015] showed how to instead first learn an approximation to C from the given samples and then compute cost shares for the learned function, but their results hold only for a limited number of games and cost functions C . We show that a more powerful approach is to compute cost shares directly from the data, without explicitly learning the cost function first.

3 Approximating the Core from Samples

In this section, we consider the problem of finding cost allocations from samples that satisfy relaxations of the core. A natural approach to this problem is to first learn the underlying model, C , from the data and to then compute a cost allocation for the learned function. As shown in Balcan et al. [2015], this approach works if C is PAC-learnable, but there exist functions C that are not PAC-learnable and for which a cost allocation that approximately satisfies the core can still be computed. The main result of this section shows that a cost allocation that approximates the core property can be computed from samples for *any* function with a non-empty core. We first give a sample complexity bound that is linear in the number n of players, a result which was independently discovered by Balcan et al. [2016]. With a more intricate analysis, we then improve this sample complexity to be logarithmic in n , but at the cost of a weaker relaxation. Our algorithm, which runs in polynomial time, directly computes a cost allocation that empirically satisfies the core property, i.e., it satisfies the core property on all of the samples. We argue, by leveraging VC-dimension and Rademacher complexity-based generalization bounds, that the same cost allocation will likely satisfy the core property on newly drawn samples as well. We also propose a stronger notion of the approximate core, and prove that it cannot be computed by any algorithm. This hardness result is information theoretic and is not due to running time limitations. The proofs in this section are deferred to Appendix B.

We begin by defining three notions of the approximate core: the probably approximately stable (Balcan et al. [2016]), mostly approximately stable, and probably mostly approximately stable cores.

Definition 3. Given $\delta, \epsilon > 0$, a cost allocation ψ such that $\sum_{i \in N} \psi_i = C(N)$ is in

- the **probably approximately stable** core if $\Pr_{S \sim \mathcal{D}} [\sum_{i \in S} \psi_i \leq C(S)] \geq 1 - \delta$ for all \mathcal{D} (see Balcan et al. [2015]),
- the **mostly approximately stable** core over \mathcal{D} if $(1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S)$ for all $S \subseteq N$,
- the **probably mostly approximately stable** core if $\Pr_{S \sim \mathcal{D}} [(1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S)] \geq 1 - \delta$ for all \mathcal{D} ,

For each of these notions, our goal is to *efficiently compute* a cost allocation in the approximate core, in the following sense.

Definition 4. A cost allocation ψ is *efficiently computable* for the class of functions \mathcal{C} over distribution \mathcal{D} , if for all $C \in \mathcal{C}$ and any $\Delta, \delta, \epsilon > 0$, given $C(N)$ and $m = \text{poly}(n, 1/\Delta, 1/\delta, 1/\epsilon)$ samples $(S_j, C(S_j))$ with each S_j drawn i.i.d. from distribution \mathcal{D} , there exists an algorithm that computes ψ with probability at least $1 - \Delta$ over both the samples and the choices of the algorithm.

We refer to the number of samples required to compute approximate cores as the *sample complexity* of the algorithm. We first present our result for computing a probably approximately stable core with sample complexity that is linear in the number of players, which was also independently discovered by Balcan et al. [2016].

Theorem 1. *The class of functions with a non-empty core has cost shares in the probably approximately stable core that are efficiently computable. The sample complexity is*

$$O\left(\frac{n + \log(1/\Delta)}{\delta}\right).$$

The full proof of Theorem 1 is in Appendix B, and can be summarized as follows: We define a class of halfspaces which contains the core. Since we assume that C has a non-empty core, there exists a cost allocation ψ in this class of halfspaces that satisfies both the core property on all the samples and the balance property. Given a set of samples, such a cost allocation can be computed with a simple linear program. We then use the VC-dimension of the class of halfspaces to show that the performance on the samples generalizes well to the performance on the distribution \mathcal{D} .

We next show that the sample complexity dependence on n can be improved from linear to logarithmic if we relax the goal from computing a cost allocation in the probably approximately stable core to computing one in the probably mostly approximately stable core instead. The sample complexity of our algorithm also depends on the *spread* of the function C , defined as $\frac{\max_S C(S)}{\min_{S \neq \emptyset} C(S)}$ (we assume $\min_{S \neq \emptyset} C(S) > 0$).

Theorem 2. *The class of functions with a non-empty core has cost allocations in the probably mostly approximately stable core that are efficiently computable with sample complexity*

$$\left(\frac{1-\epsilon}{\epsilon\delta}\right)^2 (128\tau(C)^2 \log(2n) + 8\tau(C)^2 \log(2/\Delta)) = O\left(\left(\frac{\tau(C)}{\epsilon\delta}\right)^2 (\log n + \log(1/\Delta))\right).$$

where $\tau(C) = \frac{\max_S C(S)}{\min_{S \neq \emptyset} C(S)}$ is the spread of C .

The full proof of Theorem 2 is in Appendix B. Its main steps are:

1. We find a cost allocation which satisfies the core property on all samples, restricting the search to cost allocations with bounded ℓ_1 -norm. Such a cost allocation can be found efficiently since the space of such cost allocations is convex.
2. The analysis begins by bounding the ℓ_1 -norm of any vector in the core (Lemma 3). Combined with the assumption that the core is non-empty, this implies that a cost allocation ψ satisfying the previous conditions exists.
3. Let $[x]_+$ denote the function $x \mapsto \max(x, 0)$. Consider the following “loss” function:

$$\left[\frac{\sum_{i \in S} \psi_i}{C(S)} - 1\right]_+$$

This loss function is convenient since it is equal to 0 if and only if the core property is satisfied for S and it is 1-Lipschitz, which is used in the next step.

4. Next, we bound the difference between the empirical loss and the expected loss for all ψ with a known result using the Rademacher complexity of linear predictors with low ℓ_1 norm over ρ -Lipschitz loss functions (Theorem 10).
5. Finally, given ψ which approximately satisfies the core property in expectation, we show that ψ is in the probably mostly approximately stable core by Markov’s inequality (Lemma 4).

Since we obtained a probably mostly approximately stable core, a natural question is if it is possible to compute cost allocations that are mostly approximately stable over natural distributions. The answer is negative in general: even for the restricted class of monotone submodular functions, which always have a solution in the core, the core cannot be mostly approximated from samples, even over the uniform distribution. The full proof of this impossibility theorem is in Appendix B.

Theorem 3. *Cost allocations ψ in the $(1/2 + \epsilon)$ -mostly approximately stable core, i.e., such that for all S ,*

$$\left(\frac{1}{2} + \epsilon\right) \cdot \sum_{i \in S} \psi_i \leq C(S),$$

cannot be computed for monotone submodular functions over the uniform distribution, for any constant $\epsilon > 0$.

4 Approximating the Shapley Value from Samples

We turn our attention to the STATISTICAL COST SHARING problem in the context of the Shapley value. Since the Shapley value exists and is unique for all functions, a natural relaxation is to simply approximate this value from samples. The distributions we consider in this section are the uniform distribution, and more generally product distributions, which are the standard distributions studied in the learning literature for combinatorial functions Balcan and Harvey [2011], Balcan et al. [2012], Feldman and Kothari [2014], Feldman and Vondrak [2014]. It is easy to see that we need some restrictions on the distribution \mathcal{D} (for example, if the empty set is drawn with probability one, the Shapley value cannot be approximated).

For submodular functions with bounded curvature, we prove approximation bounds when samples are drawn from the uniform or a bounded product distribution, and also show that the bound for the uniform distribution is tight. However, we show that the Shapley value cannot be approximated from samples even for coverage functions (which are a special case of submodular functions) and the uniform distribution. Since coverage functions are learnable from samples, this implies the counter-intuitive observation that learnability does not imply that the Shapley value is approximable from samples. We defer the full proofs to Appendix C.

Definition 5. An algorithm α -approximates, $\alpha \in (0, 1]$, the Shapley value of cost functions \mathcal{C} over distribution \mathcal{D} , if, for all $C \in \mathcal{C}$ and all $\delta > 0$, given $\text{poly}(n, 1/\delta, 1/(1-\alpha))$ samples from \mathcal{D} , it computes Shapley value estimates $\tilde{\phi}_C$ such that $\alpha\phi_i \leq \tilde{\phi}_i \leq \frac{1}{\alpha}\phi_i$ for all $i \in N$ such that $\phi_i \geq 1/\text{poly}(n)$ ¹ with probability at least $1 - \delta$ over both the samples and the choices made by the algorithm.

We consider submodular functions with bounded curvature, a common assumption in the submodular maximization literature Iyer and Bilmes [2013], Iyer et al. [2013], Sviridenko et al. [2015], Vondrák [2010]. Intuitively, the curvature of a submodular function bounds by how much the marginal contribution of an element can decrease. This property is useful since the Shapley value of an element can be written as a weighted sum of its marginal contributions over all sets.

Definition 6. A monotone submodular function C has curvature $\kappa \in [0, 1]$ if $C_{N \setminus \{i\}}(i) \geq (1 - \kappa)C(i)$ for all $i \in N$. This curvature is bounded if $\kappa < 1$.

An immediate consequence of this definition is that $C_S(i) \geq (1 - \kappa)C_T(i)$ for all S, T such that $i \notin S \cup T$ by monotonicity and submodularity. The main tool used is estimates \tilde{v}_i of expected marginal contributions $v_i = \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C_S(i)]$ where $\tilde{v}_i = \text{avg}(\mathcal{S}_i) - \text{avg}(\mathcal{S}_{-i})$ is the difference between the average value of samples containing i and the average value of samples not containing i .

Theorem 4. Monotone submodular functions with bounded curvature κ have Shapley value that is $\sqrt{1 - \kappa} - \epsilon$ approximable from samples over the uniform distribution, which is tight, and $1 - \kappa - \epsilon$ approximable over any bounded product distribution for any constant $\epsilon > 0$.

Consider the algorithm which computes $\tilde{\phi}_i = \tilde{v}_i$. Note that $\phi_i = \mathbb{E}_\sigma [C(A_{\sigma < i} \cup \{i\}) - C(A_{\sigma < i})] \geq (1 - \kappa)v_i > \frac{1-\kappa}{1+\epsilon}\tilde{v}_i > (1 - \kappa - \epsilon)\tilde{v}_i$ where the first inequality is by curvature and the second by Lemma 5 which shows that the estimates \tilde{v}_i of v_i are arbitrarily good. The other direction follows similarly. The $\sqrt{1 - \kappa}$ result is the main technical component of the upper bound. We describe two main steps:

1. The expected marginal contribution $\mathbb{E}_{S \sim \mathcal{U} | i \notin S, |S|=j} [C_S(i)]$ of i to a uniformly random set S of size j is decreasing in j , which is by submodularity.
2. Since a uniformly random set has size concentrated close to $n/2$, this implies that roughly half of the terms in the summation $\phi_i = (\sum_{j=0}^{n-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)]) / n$ are greater than v_i and the other half of the terms are smaller.

For the tight lower bound, we show that there exists two functions that cannot be distinguished from samples w.h.p. and that have an element with Shapley value which differs by an α^2 factor.

We show that the Shapley value of coverage (and submodular) functions are not approximable from samples in general, even though coverage functions are PMAC-learnable (Balcan and Harvey [2011]) from samples over any distribution Badanidiyuru et al. [2012].

Theorem 5. There exists no constant $\alpha > 0$ such that coverage functions have Shapley value that is α -approximable from samples over the uniform distribution.

5 Data Dependent Shapley Value

The general impossibility result for computing the Shapley value from samples arises from the fact that the concept was geared towards the query model, where the algorithm can ask for the cost of any set $S \subseteq N$. In this section, we develop an analogue that is distribution-dependent. We denote it by $\phi^{C, \mathcal{D}}$ with respect to both C and \mathcal{D} . We define four natural distribution-dependent axioms resembling

¹See Appendix C for general definition.

the Shapley value axioms, and then prove that our proposed value is the unique solution satisfying them. This value can be approximated arbitrarily well in the statistical model for all functions. The proofs are deferred to Appendix D. We start by stating the four axioms.

Definition 7. *The data-dependent axioms for cost sharing functions are:*

- **Balance:** $\sum_{i \in N} \phi_i^{\mathcal{D}} = \mathbb{E}_{S \sim \mathcal{D}}[C(S)],$
- **Symmetry:** *for all i and j , if $\Pr_{S \sim \mathcal{D}}[|S \cap \{i, j\}| = 1] = 0$ then $\phi_i^{\mathcal{D}} = \phi_j^{\mathcal{D}},$*
- **Zero element:** *for all i , if $\Pr_{S \sim \mathcal{D}}[i \in S] = 0$ then $\phi_i^{\mathcal{D}} = 0,$*
- **Additivity:** *for all i , if $\mathcal{D}_1, \mathcal{D}_2, \alpha,$ and β such that $\alpha + \beta = 1,$ $\phi_i^{\alpha \mathcal{D}_1 + \beta \mathcal{D}_2} = \alpha \phi_i^{\mathcal{D}_1} + \beta \phi_i^{\mathcal{D}_2}$ where $\Pr[S \sim \alpha \mathcal{D}_1 + \beta \mathcal{D}_2] = \alpha \cdot \Pr[S \sim \mathcal{D}_1] + \beta \cdot \Pr[S \sim \mathcal{D}_2].$*

The similarity to the original Shapley value axioms is readily apparent. The main distinction is that we expect these to hold with regard to \mathcal{D} , which captures the frequency with which different coalitions S occur. Interpreting the axioms one by one, the balance property ensures that the expected cost is always accounted for. The symmetry axiom states that if two elements always occur together, they should have the same share, since they are indistinguishable. If an element is never observed, then it should have zero share. Finally costs should combine in a linear manner according to the distribution.

The data-dependent Shapley value is

$$\phi_i^{\mathcal{D}} := \sum_{S: i \in S} \Pr[S \sim \mathcal{D}] \cdot \frac{C(S)}{|S|}.$$

Informally, for all set S , the cost $C(S)$ is divided equally between all elements in S and is weighted with the probability that S occurs according to \mathcal{D} . The main appeal of this cost allocation is the following theorem.

Theorem 6. *The data-dependent Shapley value is the unique value satisfying the four data-dependent axioms.*

The data-dependent Shapley value can be approximated from samples with the following empirical data-dependent Shapley value:

$$\tilde{\phi}_i^{\mathcal{D}} = \frac{1}{m} \sum_{S_j: i \in S_j} \frac{C(S_j)}{|S_j|}.$$

These estimates are arbitrarily good with arbitrarily high probability.

Theorem 7. *The empirical data-dependent Shapley value approximates the data-dependent Shapley value arbitrarily well, i.e.,*

$$|\tilde{\phi}_i^{\mathcal{D}} - \phi_i^{\mathcal{D}}| < \epsilon$$

with $\text{poly}(n, 1/\epsilon, 1/\delta)$ samples and with probability at least $1 - \delta$ for any $\delta, \epsilon > 0.$

6 Discussion and Future Work

We follow a recent line of work that studies classical algorithmic problems from a statistical perspective, where the input is restricted to a collection of samples. Our results fall into two categories, we give results for approximating the Shapley value and the core and propose new cost sharing concepts that are tailored for the statistical framework. We use techniques from multiple fields that encompass statistical machine learning, combinatorial optimization, and, of course, cost sharing. The cost sharing literature being very rich, the number of directions for future work are considerable. Obvious avenues include studying other cost sharing methods in this statistical framework, considering other classes of functions to approximate known methods, and improving the sample complexity of previous algorithms. More conceptually, an exciting modeling question arises when designing "desirable" axioms from data. Traditionally these axioms only depended on the cost function, whereas in this model they can depend on both the cost function and the distribution, providing an interesting interplay.

References

- Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- Yoram Bachrach, Evangelos Markakis, Ezra Resnick, Ariel D Procaccia, Jeffrey S Rosenschein, and Amin Saberi. Approximating power indices: theoretical and empirical analysis. *Autonomous Agents and Multi-Agent Systems*, 20(2):105–122, 2010.
- Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu, Robert Kleinberg, Noam Nisan, and Tim Roughgarden. Sketching valuation functions. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1025–1035. Society for Industrial and Applied Mathematics, 2012.
- Maria-Florina Balcan and Nicholas JA Harvey. Learning submodular functions. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 793–802. ACM, 2011.
- Maria-Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. Learning valuation functions. In *COLT*, volume 23, pages 4–1, 2012.
- Maria-Florina Balcan, Ariel D. Procaccia, and Yair Zick. Learning cooperative games. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 475–481, 2015.
- Maria-Florina Balcan, Ariel D Procaccia, and Yair Zick. Learning cooperative games. *arXiv preprint arXiv:1505.00039v2*, 2016.
- Eric Balkanski and Yaron Singer. Mechanisms for fair attribution. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 529–546. ACM, 2015.
- Olga N Bondareva. Some applications of linear programming methods to the theory of cooperative games. *Problemy kibernetiki*, 10:119–139, 1963.
- Xiaotie Deng, Toshihide Ibaraki, and Hiroshi Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Mathematics of Operations Research*, 24(3):751–766, 1999.
- Nikhil R Devanur, Milena Mihail, and Vijay V Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39(1):11–22, 2005.
- Shaheen S Fatima, Michael Wooldridge, and Nicholas R Jennings. A linear approximation method for the shapley value. *Artificial Intelligence*, 172(14):1673–1699, 2008.
- Joan Feigenbaum, Christos Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions (preliminary version). In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 218–227. ACM, 2000.
- Vitaly Feldman and Pravesh Kothari. Learning coverage functions and private release of marginals. In *COLT*, pages 679–702, 2014.
- Vitaly Feldman and Jan Vondrak. Optimal bounds on approximation of submodular and xos functions by juntas. In *Information Theory and Applications Workshop (ITA), 2014*, pages 1–10. IEEE, 2014.
- Donald B Gillies. Solutions to general non-zero-sum games. *Contributions to the Theory of Games*, 4(40):47–85, 1959.
- Michel X Goemans and Martin Skutella. Cooperative facility location games. *Journal of Algorithms*, 50(2):194–214, 2004.
- Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2014. URL <http://arxiv.org/abs/1412.6572>.
- Nicole Immorlica, Mohammad Mahdian, and Vahab S Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms (TALG)*, 4(2):24, 2008.

- Rishabh K Iyer and Jeff A Bilmes. Submodular optimization with submodular cover and submodular knapsack constraints. In *Advances in Neural Information Processing Systems*, pages 2436–2444, 2013.
- Rishabh K Iyer, Stefanie Jegelka, and Jeff A Bilmes. Curvature and optimal algorithms for learning and minimizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 2742–2750, 2013.
- Kamal Jain and Vijay V Vazirani. Equitable cost allocations via primal-dual-type algorithms. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 313–321. ACM, 2002.
- David Liben-Nowell, Alexa Sharp, Tom Wexler, and Kevin Woods. Computing shapley value in supermodular coalitional games. In *International Computing and Combinatorics Conference*, pages 568–579. Springer, 2012.
- Irwin Mann. *Values of large games, IV: Evaluating the electoral college by Montecarlo techniques*. Rand Corporation, 1960.
- Hervé Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
- Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
- Martin Pál and Éva Tardos. Group strategy proof mechanisms via primal-dual algorithms. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 584–593. IEEE, 2003.
- Alvin E Roth. *The Shapley value: essays in honor of Lloyd S. Shapley*. Cambridge University Press, 1988.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. 2014.
- Lloyd S Shapley. On balanced sets and cores. *Naval research logistics quarterly*, 14(4):453–460, 1967.
- LS Shapley. A value for n-person games¹. 1953.
- Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1134–1148. SIAM, 2015.
- Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *CoRR*, abs/1312.6199, 2013. URL <http://arxiv.org/abs/1312.6199>.
- John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. 1944.
- Jan Vondrák. Submodularity and curvature: the optimal algorithm. *RIMS Kokyuroku Bessatsu B*, 23: 253–266, 2010.
- Eyal Winter. The shapley value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.

Appendix

A Concentration Bounds

Lemma 1 (Chernoff Bound). *Let X_1, \dots, X_n be independent indicator random variables with values in $\{0, 1\}$. Let $X = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[X]$. For $0 < \delta < 1$,*

$$\Pr[X \leq (1 - \delta)\mu] \leq e^{-\mu\delta^2/2} \quad \text{and} \quad \Pr[X \geq (1 + \delta)\mu] \leq e^{-\mu\delta^2/3}.$$

Lemma 2 (Hoeffding's inequality). *Let X_1, \dots, X_n be independent random variables with values in $[0, b]$. Let $X = \frac{1}{m} \sum_{i=1}^m X_i$ and $\mu = \mathbb{E}[X]$. Then for every $0 < \epsilon < 1$,*

$$\Pr[|X - \mathbb{E}[X]| \geq \epsilon] \leq 2e^{-2m\epsilon^2/b^2}.$$

B Missing Definitions and Analysis from Section 3

We first show the result with linear sample complexity, then the result with logarithmic sample complexity, and finally the impossibility result.

B.1 Linear sample complexity for probably approximately stable core

We first state the generalization error obtained for a class of functions with VC-dimension d .

Theorem 8 (Shalev-Shwartz and Ben-David [2014], Theorem 6.8). *Let \mathcal{H} be a hypothesis class of functions from a domain \mathcal{X} to $\{-1, 1\}$ and $f : \mathcal{X} \mapsto \{-1, 1\}$ be some "correct" function. Assume that \mathcal{H} has VC-dimension d . Then, there is an absolute constant c such that with $m \geq c(d + \log(1/\Delta))/\delta$ i.i.d. samples $\mathbf{x}^1, \dots, \mathbf{x}^m \sim \mathcal{D}$,*

$$\left| \Pr_{\mathbf{x} \sim \mathcal{D}} [h(\mathbf{x}) \neq f(\mathbf{x})] - \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{h(\mathbf{x}^i) \neq f(\mathbf{x}^i)} \right| \leq \delta$$

for all $h \in \mathcal{H}$, with probability $1 - \Delta$ over the samples.

We use a special case of the class of halfspaces, for which we know the VC-dimension.

Theorem 9 (Shalev-Shwartz and Ben-David [2014], Theorem 9.2). *The class of functions $\{\mathbf{x} \mapsto \text{sign}(\mathbf{w}^\top \mathbf{x}) : \mathbf{w} \in \mathbb{R}^n\}$ has VC-dimension n .*

We first define a class of functions that contains the core, and prove that it has low VC-dimension. Given a sample S , define \mathbf{x}^S such that $x_i^S = \mathbb{1}_{i \in S}$ for $i \in [n]$ and $x_{n+1}^S = C(S)$. Note that $\text{sign}(\sum_{i=1}^n \psi_i x_i^S - x_{n+1}^S) = -1$ if the core property is satisfied for sample S . We now bound the VC-dimension of this hypothesis class of functions induced by cost allocations ψ .

Corollary 1. *The class of functions $\mathcal{H}^{\text{core}} = \{\mathbf{x} \mapsto \text{sign}(\sum_{i=1}^n \psi_i x_i - x_{n+1}) : \psi \in \mathbb{R}^n, \sum_i \psi_i = C(N)\}$ has VC-dimension at most $n + 1$.*

Proof. We combine the observation that $\{\mathbf{x} \mapsto \text{sign}(\sum_{i=1}^n w_i x_i - x_{n+1}) : \mathbf{w} \in \mathbb{R}^n, \sum_i w_i = C(N)\} \subseteq \{\mathbf{x} \mapsto \text{sign}(\mathbf{w}^\top \mathbf{x}) : \mathbf{w} \in \mathbb{R}^{n+1}\}$ with the well-known fact that the VC-dimension of \mathcal{H}' is at most the VC-dimension of \mathcal{H} for $\mathcal{H}' \subseteq \mathcal{H}$. \square

It remains to show how to optimize over functions in this class.

Theorem 1. *The class of functions with a non-empty core has cost shares in the probably approximately stable core that are efficiently computable. The sample complexity is*

$$O\left(\frac{n + \log(1/\Delta)}{\delta}\right).$$

Proof. Let ψ be a cost allocation which satisfies both the core property on all the samples and the balance property, i.e., $\sum_{i \in S} \psi_i \leq C(S)$ for all samples S and $\sum_{i \in N} \psi_i = C(N)$. Note that such a

cost allocation exists since we assume that C has a non-empty core. Given the set of samples, it can be computed with a simple linear program. We argue that ψ is probably approximately stable.

Define $h(x) = \text{sign}(\sum_{i=1}^n \psi_i x_i^S - x_{n+1}^S)$ and $f(\mathbf{x}) = -1$ for all \mathbf{x} . Since the core property is satisfied on all the samples, $\frac{1}{m} \sum_{i=1}^m \mathbb{1}_{h(\mathbf{x}) \neq f(\mathbf{x})} = 0$. Thus, by Theorem 8,

$$\begin{aligned} \Pr_{S \sim \mathcal{D}} \left[\sum_{i \in S} \psi_i \leq C(S) \right] &= 1 - \Pr_{\mathbf{x}^S: S \sim \mathcal{D}} \left[\text{sign} \left(\sum_{i=1}^n \psi_i x_i^S - x_{n+1}^S \right) \neq -1 \right] \\ &= 1 - \Pr_{\mathbf{x}^S: S \sim \mathcal{D}} [h(\mathbf{x}^S) \neq f(\mathbf{x}^S)] \\ &= 1 - \left| \Pr_{\mathbf{x}^S: S \sim \mathcal{D}} [h(\mathbf{x}^S) \neq f(\mathbf{x}^S)] - \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{h(\mathbf{x}^S) \neq f(\mathbf{x}^S)} \right| \\ &\geq 1 - \delta \end{aligned}$$

with $O((n + \log(1/\Delta))/\delta)$ samples. \square

B.2 Logarithmic sample complexity for probably mostly approximately stable cores

The following result follows from the Rademacher complexity of linear classes.

Theorem 10 (Shalev-Shwartz and Ben-David [2014], Theorem 26.15). *Suppose that \mathcal{D} is a distribution over $\mathcal{X} \times \mathbb{R}$ such that with probability 1 we have that $\|\mathbf{x}\|_\infty \leq R$. Let $\mathcal{H} = \{\mathbf{w} \in \mathbb{R}^d : \|\mathbf{w}\|_1 \leq B\}$ and let $\ell : \mathcal{H} \times (\mathcal{X} \times \mathbb{R}) \rightarrow \mathbb{R}$ be a loss function of the form*

$$\ell(\mathbf{w}, (\mathbf{x}, y)) = \phi(\mathbf{w}^\top \mathbf{x}, y)$$

such that for all $y \in \mathbb{R}$, $a \mapsto \phi(a, y)$ is an ρ -Lipschitz function and such that $\max_{a \in [-BR, BR]} |\phi(a, y)| \leq c$. Then, for all $\mathbf{w} \in \mathcal{H}$ and any $\Delta \in (0, 1)$, with probability of at least $1 - \Delta$ over m i.i.d. samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_m, y_m)$ from \mathcal{D} ,

$$\mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}} [\ell(\mathbf{w}, (\mathbf{x}, y))] \leq \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{w}, (\mathbf{x}_i, y_i)) + 2\rho BR \sqrt{\frac{2 \log(2d)}{m}} + c \sqrt{\frac{2 \log(2/\Delta)}{m}}.$$

We first bound the ℓ_1 norm of vectors in the core to bound the space of linear functions that we search over.

Lemma 3. *Assume that ψ is a vector in the core, then $\|\psi\|_1 \leq 2 \max_S |C(S)|$.*

Proof. Fix some vector ψ in the core. Let A be the set of elements i such that $\psi_i \geq 0$ and B be the remaining elements. First note that

$$\sum_{i \in A} \psi_i \leq C(A) \leq \max_S |C(S)|$$

where the first inequality is by the core property. Next, note that

$$0 \leq C(N) = \sum_{i \in A} \psi_i + \sum_{i \in B} \psi_i \leq \max_S |C(S)| + \sum_{i \in B} \psi_i$$

where the equality is by the balance property, so $\sum_{i \in B} \psi_i \geq -\max_S |C(S)|$. Thus,

$$\|\psi\|_1 = \sum_{i \in A} \psi_i - \sum_{i \in B} \psi_i \leq \max_S |C(S)| + \max_S |C(S)|.$$

\square

We can thus focus on bounded cost allocations $\psi \in \mathcal{H}$ where

$$\mathcal{H} := \left\{ \psi : \psi \in \mathbb{R}^n, \|\psi\|_1 \leq 2 \max_S |C(S)| \right\}.$$

The next lemma shows that if the core property approximately holds in expectation, then it is likely to approximately hold.

Lemma 4. For any $0 < \epsilon, \delta < 1$ and cost allocation ψ ,

$$\mathbb{E}_{S \sim \mathcal{D}} \left[\left[\frac{\sum_{i \in S} \psi_i}{C(S)} - 1 \right]_+ \right] \leq \frac{\epsilon \delta}{1 - \epsilon} \Rightarrow \Pr_{S \sim \mathcal{D}} \left[(1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S) \right] \geq 1 - \delta.$$

Proof. For any $a > 0$ and nonnegative random variable X , by Markov's inequality we have $\Pr[X \leq a] \geq 1 - E[X]/a$. By letting $a = \epsilon/(1 - \epsilon)$, $X = [(\sum_{i \in S} \psi_i)/C(S) - 1]_+$, and observing that

$$\left[\frac{\sum_{i \in S} \psi_i}{C(S)} - 1 \right]_+ \leq \frac{\epsilon}{1 - \epsilon} \Rightarrow \frac{\sum_{i \in S} \psi_i}{C(S)} - 1 \leq \frac{\epsilon}{1 - \epsilon} \Rightarrow (1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S),$$

we obtain $\Pr_{S \sim \mathcal{D}} [(1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S)] \geq 1 - \delta$. \square

Combining Theorem 10, Lemma 3, and Lemma 4, we obtain the main result.

Theorem 2. The class of functions with a non-empty core has cost allocations in the probably mostly approximately stable core that are efficiently computable with sample complexity

$$\left(\frac{1 - \epsilon}{\epsilon \delta} \right)^2 (128\tau(C)^2 \log(2n) + 8\tau(C)^2 \log(2/\Delta)) = O \left(\left(\frac{\tau(C)}{\epsilon \delta} \right)^2 (\log n + \log(1/\Delta)) \right).$$

where $\tau(C) = \frac{\max_S C(S)}{\min_{S \neq \emptyset} C(S)}$ is the spread of C .

Proof. Fix $C \in \mathcal{C}$. Suppose we are given m samples from \mathcal{D} . We pick $\psi^* \in \mathcal{H}$ such that core property holds on all the samples and such that the balance property holds ($\sum_{i \in N} \psi_i = C(N)$). This cost allocation ψ^* can be found efficiently since the collection of such ψ is a convex set. By the assumption that C has at least one vector in the core and by Lemma 3, such a ψ^* exists. Given $S \sim \mathcal{D}$, define \mathbf{x}^S such that $x_i^S = \mathbb{1}_{i \in S}/C(S)$. Fix $y = 1$. Define the loss function ℓ as follows,

$$\ell(\psi, (\mathbf{x}^S, y)) := [\psi^\top \mathbf{x}^S - y]_+ = \left[\frac{\sum_{i \in S} \psi_i}{C(S)} - 1 \right]_+$$

We wish to use Theorem 10 with $R = 1/\min_{S \neq \emptyset} |C(S)|$, $B = 2 \max_S |C(S)|$, $\phi(a, y) = [a - y]_+$, $\rho = 1$, and $c = \tau(C)$. We verify that all the conditions hold. First note that without loss of generality, samples where $S = \emptyset$ can be ignored, so $\|\mathbf{x}^S\|_\infty \leq 1/\min_{S \neq \emptyset} |C(S)|$ for all S . Next, $\|\psi\|_1 \leq 2 \max_S |C(S)|$ for $\psi \in \mathcal{H}$ by definition of \mathcal{H} . The loss function ℓ is of the form $\ell(\psi, (\mathbf{x}, y)) = \phi(\psi^\top \mathbf{x}, y) = [\psi^\top \mathbf{x} - y]_+$ such that $a \mapsto \phi(a, y) = [a - y]_+$ is a 1-Lipschitz function and such that $\max_{a \in [-BR, BR]} |\phi(a, y)| \leq 2 \max_S |C(S)|/\min_{S \neq \emptyset} |C(S)| = 2\tau(C)$. In addition, note that

$$\frac{1}{m} \sum_{i=1}^m \ell(\psi^*, (\mathbf{x}^{S_i}, 1)) = 0$$

since the core property holds on all the samples. Thus, by Theorem 10,

$$\mathbb{E}_{S \sim \mathcal{D}} \left[\left[\frac{\sum_{i \in S} \psi_i^*}{C(S)} - 1 \right]_+ \right] = \mathbb{E}_{\mathbf{x}^S: S \sim \mathcal{D}} [\ell(\psi^*, (\mathbf{x}^S, 1))] \leq 4\tau(C) \sqrt{\frac{2 \log(2n)}{m}} + \tau(C) \sqrt{\frac{2 \log(2/\Delta)}{m}}.$$

Choose any $\epsilon, \delta > 0$. If the number of samples m is chosen as in the statement of the theorem, then the righthand side of the above inequality will be less than $\frac{\epsilon \delta}{1 - \epsilon}$. Thus by Lemma 4,

$$\Pr_{S \sim \mathcal{D}} \left[(1 - \epsilon) \sum_{i \in S} \psi_i \leq C(S) \right] \geq 1 - \delta,$$

which completes the proof. \square

B.3 Hardness of mostly approximately stable core

We now give the missing analysis for the impossibility result of approximating the core.

Theorem 3. *Cost allocations ψ in the $(1/2 + \epsilon)$ -mostly approximately stable core, i.e., such that for all S ,*

$$\left(\frac{1}{2} + \epsilon\right) \cdot \sum_{i \in S} \psi_i \leq C(S),$$

cannot be computed for monotone submodular functions over the uniform distribution, for any constant $\epsilon > 0$.

Proof. The ground set of elements is partitioned in ϵ^{-1} sets $A_1, \dots, A_{\epsilon^{-1}}$ of size ϵn for some small constant $\epsilon > 0$. Let $\mathcal{C} = \{C^{A_i} : i \in [\epsilon^{-1}]\}$ where

$$C^A(S) = |(N \setminus A) \cap S| + \min\left(|A \cap S|, (1 + \epsilon) \frac{\epsilon n}{2}\right).$$

The expected number of elements of A_i in a sample S from the uniform distribution is $|A_i|/2 = \epsilon n/2$, so by the Chernoff bound

$$\Pr\left[|A_i \cap S| \geq (1 + \epsilon) \frac{\epsilon n}{2}\right] \leq e^{-\frac{\epsilon^2 n}{6}},$$

Thus, by a union bound, $C^{A_i}(S) = |S|$ over all i and all samples S with probability $1 - O(e^{-n})$ and we henceforth assume this is the case. It is therefore impossible to learn any information about the partition $A_1, \dots, A_{\epsilon^{-1}}$ from samples. Any cost allocation ψ computed by an algorithm given samples from C^{A_i} is thus *independent* of i .

Next, consider such a cost allocation ψ independent of i satisfying the balance property. There exists A_i such that $\sum_{j \in A_i} \psi_j > (1 - \epsilon)n\epsilon$ since $\sum_{j \in N} \psi_j = C(N) > (1 - \epsilon)n$ by the balance property. In addition, $C^{A_i}(A_i) = (1 + \epsilon)\epsilon n/2$. We obtain

$$\sum_{j \in A_i} \psi_j > (1 - \epsilon)n\epsilon = \frac{(1 - \epsilon)}{(1 + \epsilon)} 2C^{A_i}(A_i).$$

Thus, the core property is violated by a $1/2 + \epsilon'$ factor for set A_i and function C^{A_i} , and for any constant $\epsilon' > 0$ by picking ϵ sufficiently small. \square

C Missing Analysis from Section 4

We give a complete definition of approximate Shapley values.

Definition 8. *An algorithm α -approximates, $\alpha \in (0, 1]$, the Shapley value of a family of cost functions \mathcal{C} over distribution \mathcal{D} , if, for all $C \in \mathcal{C}$ and all $\delta > 0$, given $\text{poly}(n, 1/\delta, 1/(1-\alpha))$ samples from \mathcal{D} , it computes Shapley value estimates $\tilde{\phi}_C$ such that for*

- *positive bounded Shapley value, if $\phi_i \geq 1/\text{poly}(n)$, then $\alpha\phi_i \leq \tilde{\phi}_i \leq \frac{1}{\alpha}\phi_i$;*
- *negative bounded Shapley value, if $\phi_i \leq -1/\text{poly}(n)$, then $\frac{1}{\alpha}\phi_i \leq \tilde{\phi}_i \leq \alpha\phi_i$;*
- *small Shapley value, if $|\phi_i| < 1/\text{poly}(n)$, then $|\phi_i - \tilde{\phi}_i| = o(1)$.*

for all $i \in N$ with probability at least $1 - \delta$ over both the samples and the choices made by the algorithm.

Next, we show that the estimates \tilde{v}_i of the expected marginal contribution $v_i = \mathbb{E}_{S \sim \mathcal{D} | i \notin S}[C_S(i)]$ are arbitrarily good. Recall that \mathcal{S}_i and \mathcal{S}_{-i} are the collections of samples containing element i and not containing it respectively and that $\text{avg}(\mathcal{S}) = (\sum_{S \in \mathcal{S}} C(S))/|\mathcal{S}|$ is the average value of the samples in \mathcal{S} . Let $v_i = C(i)$ and $\tilde{v}_i = \text{avg}(\mathcal{S}_i) - \text{avg}(\mathcal{S}_{-i})$.

Lemma 5. *The expected marginal contribution of an element i to a random set from a bounded product distribution \mathcal{D} not containing i is estimated arbitrarily well by \tilde{v}_i , i.e., for all $i \in N$ and given $\text{poly}(n, 1/\delta, 1/\epsilon)$ samples,*

$$\begin{aligned} (1 - \epsilon)v_i \leq \tilde{v}_i \leq (1 + \epsilon)v_i & \quad \text{if } v_i \geq 1/\text{poly}(n) \\ |v_i - \tilde{v}_i| \leq \epsilon & \quad \text{if } |v_i| < 1/\text{poly}(n) \\ (1 + \epsilon)v_i \leq \tilde{v}_i \leq (1 - \epsilon)v_i & \quad \text{if } v_i \leq -1/\text{poly}(n) \end{aligned}$$

with probability at least $1 - \delta$ for any $\delta > 0$.

Proof. Note that

$$v_i = \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C_S(i)] = \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C(S \cup i)] - \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C(S)] = \mathbb{E}_{S \sim \mathcal{D} | i \in S} [C(S)] - \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C(S)].$$

where the second equality is since \mathcal{D} is a product distribution. We also have that $\mathbb{E}[\text{avg}(\mathcal{S}_i)] = \mathbb{E}_{S \sim \mathcal{D} | i \in S} [C(S)]$ and $\mathbb{E}[\text{avg}(\mathcal{S}_{-i})] = \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C(S)]$. Since marginal probabilities of the product distributions are assumed to be bounded from below and above by $1/\text{poly}(n)$ and $1 - 1/\text{poly}(n)$ respectively, $|\mathcal{S}_i| = m/\text{poly}(n)$ and $|\mathcal{S}_{-i}| = m/\text{poly}(n)$ for all i by Chernoff bound. In addition, $\max_S C(S)$ is assumed to be bounded by $\text{poly}(n)$. So by Hoeffding's inequality,

$$\Pr \left(\left| \text{avg}(\mathcal{S}_i) - \mathbb{E}_{S \sim \mathcal{D} | i \in S} [C(S)] \right| \geq |v_i|\epsilon/2 \right) \leq 2e^{-\frac{m(|v_i|\epsilon)^2}{\text{poly}(n)}},$$

for $0 < \epsilon < 2/v_i$ and

$$\Pr \left(\left| \text{avg}(\mathcal{S}_{-i}) - \mathbb{E}_{S \sim \mathcal{D} | i \notin S} [C(S)] \right| \geq |v_i|\epsilon/2 \right) \leq 2e^{-\frac{m(|v_i|\epsilon)^2}{\text{poly}(n)}}.$$

Thus,

$$\Pr(|\tilde{v}_i - v_i| \geq |v_i|\epsilon) \leq 2e^{-\frac{m(|v_i|\epsilon)^2}{\text{poly}(n)}}$$

and, either $(1 - \epsilon)v_i \leq \tilde{v}_i \leq (1 + \epsilon)v_i$ if $v_i > 0$ or $(1 + \epsilon)v_i \leq \tilde{v}_i \leq (1 - \epsilon)v_i$ if $v_i < 0$, with probability at least $1 - 2e^{-\frac{m(|v_i|\epsilon)^2}{\text{poly}(n)}}$. If $|v_i| < 1/\text{poly}(n)$, we obtain $|v_i - \tilde{v}_i| < \epsilon$ with a similar analysis without any assumption on v_i . Otherwise, the bounds on the estimation hold with probability at least $1 - 2e^{-\frac{m\epsilon^2}{\text{poly}(n)}}$. \square

Theorem 4. *Monotone submodular functions with bounded curvature κ have Shapley value that is $\sqrt{1 - \kappa} - \epsilon$ approximable from samples over the uniform distribution, which is tight, and $1 - \kappa - \epsilon$ approximable over any bounded product distribution for any constant $\epsilon > 0$.*

It remains to improve the bound from $1 - \kappa - \epsilon$ to $\sqrt{1 - \kappa} - \epsilon$ for the uniform distribution. Let

$$\tilde{\phi}_i = \frac{2 - \kappa}{2\sqrt{1 - \kappa}} \cdot \tilde{v}_i$$

be the estimated Shapley value. Denote by \mathcal{U}_j the uniform distribution over all sets of size j , so

$$\phi_i = \mathbb{E}_{\sigma} [C_{A_{\sigma < i}}(i)] = \frac{1}{n} \sum_{j=0}^{n-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)].$$

The main idea to improve the loss from $1 - \kappa$ to $\sqrt{1 - \kappa}$ is to observe that v_i can be a factor $1 - \kappa$ away from the contribution $\mathbb{E}_{S \sim \mathcal{U}_{j_l} | i \notin S} [C_S(i)]$ of j to sets of low sizes $j_l \leq L := (1 - \epsilon') \cdot n/2$ or $1 - \kappa$ away from its contribution $\mathbb{E}_{S \sim \mathcal{U}_{j_h} | i \notin S} [C_S(i)]$ to sets of high sizes $j_h \geq H := (1 + \epsilon') \cdot n/2$, but not both, otherwise the curvature property would be violated as illustrated in Figure 1. The following lemma shows that $\mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)]$ is decreasing in j by submodularity.

Lemma 6. *Let C be a submodular function, then for all $j \in \{0, \dots, n - 1\}$ and all $i \in N$,*

$$\mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \geq \mathbb{E}_{S \sim \mathcal{U}_{j+1} | i \notin S} [C_S(i)]$$




Figure 1: The expected marginal contribution $E_{S \sim \mathcal{U}_j | i \notin S}[C_S(i)]$ of an element i to a set of size j as a function of j . The curvature property implies that any two points are at most a factor $1 - \kappa$ from each other. Lemma 6 shows that it is decreasing. Lemma 7 shows that the expected marginal contribution v_i of i to a uniformly random set is approximately between $E_{S \sim \mathcal{U}_L | i \notin S}[C_S(i)]$ and $E_{S \sim \mathcal{U}_H | i \notin S}[C_S(i)]$. The Shapley value of i is the average value of this expected marginal contribution over all integers $j \in \{0, \dots, n-1\}$.

Proof. By submodularity,

$$\sum_{S: |S|=j, i \notin S} C_S(i) \geq \sum_{S: |S|=j, i \notin S} \frac{1}{n-j-1} \sum_{i' \notin S \cup \{i\}} C_{S \cup \{i'\}}(i).$$

In addition, observe that by counting in two ways,

$$\sum_{S: |S|=j, i \notin S} \sum_{i' \notin S \cup \{i\}} C_{S \cup \{i'\}}(i) = (j+1) \sum_{S: |S|=j+1, i \notin S} C_S(i).$$

By combining the two previous observations,

$$\begin{aligned} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S}[C_S(i)] &= \frac{1}{|\{S : |S|=j, i \notin S\}|} \sum_{S: |S|=j, i \notin S} C_S(i) \\ &\geq \frac{1}{\binom{n-1}{j}} \frac{j+1}{n-j-1} \sum_{S: |S|=j+1, i \notin S} C_S(i) \\ &= \frac{1}{|\{S : |S|=j+1, i \notin S\}|} \sum_{S: |S|=j+1, i \notin S} C_S(i) \\ &= \mathbb{E}_{S \sim \mathcal{U}_{j+1} | i \notin S}[C_S(i)] \end{aligned}$$

□

The next lemma shows that for j slightly lower than $n/2$, the expected marginal contribution v_i of element i to a random set cannot be much larger than $\mathbb{E}_{S \sim \mathcal{U}_j | i \notin S}[C_S(i)]$, and similarly for j slightly larger than $n/2$, it cannot be much smaller.

Lemma 7. *Let C be a submodular function, then for all $i \in N$,*

$$\left(1 + \frac{e^{-\frac{\epsilon n}{6}}}{1 - \kappa}\right) \cdot \mathbb{E}_{S \sim \mathcal{U}_L | i \notin S}[C_S(i)] \geq v_i \geq (1 - e^{-\frac{\epsilon n}{6}}) \cdot \mathbb{E}_{S \sim \mathcal{U}_H | i \notin S}[C_S(i)].$$

Proof. By Chernoff bound, $L \leq |S|$ and $|S| \leq H$ with probability at least $1 - e^{-\frac{\epsilon n}{6}}$ each for S drawn from the uniform distribution. Denote the uniform distribution over all sets by \mathcal{U} . So,

$$\begin{aligned}
v_i &= \sum_{j=0}^{n-1} \Pr_{S \sim \mathcal{U} | i \notin S} [|S| = j] \cdot \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \\
&\geq \sum_{j=0}^H \Pr_{S \sim \mathcal{U} | i \notin S} [|S| = j] \cdot \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \\
&\geq \Pr_{S \sim \mathcal{U} | i \notin S} [|S| \leq H] \cdot \mathbb{E}_{S \sim \mathcal{U}_H | i \notin S} [C_S(i)] && \text{Lemma 6} \\
&\geq (1 - e^{-\frac{\epsilon n}{6}}) \cdot \mathbb{E}_{S \sim \mathcal{U}_H | i \notin S} [C_S(i)].
\end{aligned}$$

Similarly,

$$\begin{aligned}
v_i &= \sum_{j=0}^{n-1} \Pr_{S \sim \mathcal{U} | i \notin S} [|S| = j] \cdot \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \\
&\leq \Pr_{S \sim \mathcal{U} | i \notin S} [|S| < L] \cdot C(i) + \Pr_{S \sim \mathcal{U} | i \notin S} [|S| \geq L] \cdot \mathbb{E}_{S \sim \mathcal{U}_L | i \notin S} [C_S(i)] && \text{Lemma 6} \\
&\leq e^{-\frac{\epsilon n}{6}} \cdot \frac{1}{1 - \kappa} \cdot \mathbb{E}_{S \sim \mathcal{U}_L | i \notin S} [C_S(i)] + \mathbb{E}_{S \sim \mathcal{U}_L | i \notin S} [C_S(i)] && \text{curvature}
\end{aligned}$$

□

We are now ready to prove Theorem 4:

$$\begin{aligned}
\phi_i &= \frac{1}{n} \sum_{j=0}^{n-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \\
&= \frac{1}{n} \left(\sum_{j=0}^{H-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] + \sum_{i=H}^{n-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] \right) \\
&\leq \frac{1 + \epsilon'}{2} \cdot C(i) + \frac{1}{2} \cdot \mathbb{E}_{S \sim \mathcal{U}_H | i \notin S} [C_S(i)] && \text{Lemma 6} \\
&\leq \frac{1 + \epsilon'}{2(1 - \kappa)} \cdot v_i + \frac{1}{2(1 - e^{-\frac{\epsilon' n}{6}})} \cdot v_i && \text{curvature and Lemma 7} \\
&\leq \left(\frac{2 - \kappa}{2(1 - \kappa)} + c_1 \epsilon' \right) \cdot v_i \\
&\leq \left(\frac{2 - \kappa}{2(1 - \kappa)} + c_2 \epsilon' \right) \cdot \tilde{v}_i && \text{Lemma 5} \\
&= \left(\frac{1}{\sqrt{1 - \kappa} - c_3 \epsilon'} \right) \cdot \tilde{\phi}_i && \text{definition of } \tilde{\phi}_i
\end{aligned}$$

for some constants c_1, c_2, c_3 and let $\epsilon' = \epsilon/c_3$ to obtain the desired result for any ϵ . Similarly,

$$\begin{aligned}
\phi_i &= \frac{1}{n} \sum_{j=0}^L \mathbb{E}_{S \sim \mathcal{U}_j | i \notin S} [C_S(i)] + \frac{1}{n} \sum_{j=L+1}^{n-1} \mathbb{E}_{S \sim \mathcal{U}_j | i \in S} [C_S(i)] \\
&\geq \frac{1 - \epsilon'}{2} \cdot \mathbb{E}_{S \sim \mathcal{U}_L | i \notin S} [C_S(i)] + \frac{1}{2} (C(N) - C(N \setminus \{i\})) && \text{Lemma 6} \\
&\geq \frac{1 - \epsilon'}{2 \left(1 + \frac{e^{-\frac{\epsilon' n}{6}}}{1 - \kappa}\right)} \cdot v_i + \frac{1 - \kappa}{2} \cdot v_i && \text{Lemma 7 and curvature} \\
&\geq \left(\frac{2 - \kappa}{2} - c_1 \epsilon'\right) v_i \\
&\geq \left(\frac{2 - \kappa}{2} - c_2 \epsilon'\right) \tilde{v}_i && \text{Lemma 5} \\
&= (\sqrt{1 - \kappa} - c_3 \epsilon') \tilde{\phi}_i && \text{definition of } \tilde{\phi}_i
\end{aligned}$$

□

We show that this approximation is optimal. We begin with a general lemma to derive information theoretic inapproximability results for the Shapley value. This lemma shows that if there exists two functions in \mathcal{C} that cannot be distinguished from samples with high probability and that have an element with Shapley value which differs by an α^2 factor, then \mathcal{C} does not have a Shapley value that is α -approximable from samples.

Lemma 8. *Consider a family of cost functions \mathcal{C} , a constant $\alpha \in (0, 1)$, and assume there exist $C^1, C^2 \in \mathcal{C}$ such that:*

- **Indistinguishable from samples.** *With probability $1 - O(e^{-\beta n})$ over $S \sim \mathcal{D}$ for some constant $\beta > 0$,*

$$C^1(S) = C^2(S).$$

- **Gap in Shapley value.** *There exists $i \in N$ such that*

$$\phi_i^{C^1} < \alpha^2 \phi_i^{C^2}.$$

Then, \mathcal{C} does not have Shapley value that is α -approximable from samples over \mathcal{D} .

Proof. By a union bound, given m sets S_1, \dots, S_m drawn i.i.d. from \mathcal{D} with m polynomial in n , $C^1(S_j) = C^2(S_j)$ for all S_j with probability $1 - O(e^{-\beta n})$.

Let $C = C^1$ or $C = C^2$ with probability $1/2$ each. Assume the algorithm is given m samples such that $C^1(S_j) = C^2(S_j)$ and consider its (possibly randomized) choice $\tilde{\phi}_i$. Note that $\tilde{\phi}_i$ is independent of the randomization of C since C^1 and C^2 are indistinguishable to the algorithm. Since $\phi_i^{C^1} / \phi_i^{C^2} < \alpha^2$, $\tilde{\phi}_i$ is at least a factor α away from ϕ_i^C with probability at least $1/2$ over the choices of the algorithm and C . Label the cost functions so that $\tilde{\phi}_i$ is at least a factor α away from $\phi_i^{C^1}$ with probability at least $1/2$ over the choices of the algorithm. Thus, with $\delta = 1/4$, there exists no algorithm such that for all $C' \in \{C^1, C^2\}$, $\alpha \cdot \phi_i^{C'} \leq \tilde{\phi}_i^{C'} \leq \frac{1}{\alpha} \cdot \phi_i^{C'}$ with probability at least $3/4$ over both the samples and the choices of the algorithm. □

We obtain the inapproximability result by constructing two such functions.

Theorem 11. *For every $\kappa < 1$, there exists a hypothesis class of submodular functions with curvature κ that have Shapley value that is not $\sqrt{1 - \kappa} + \epsilon$ -approximable from samples over the uniform distribution, for every constant $\epsilon > 0$.*

Proof. We show that an element i^* can have Shapley value that differs by a factor $1 - \kappa$ for two functions C^1 and C^2 with curvature κ that are indistinguishable from samples. Combining with Lemma 8, we then obtain the negative result. These two functions have a simpler definition via their




Figure 2: The marginal contributions $C_S^1(i^*)$, (a), and $C_S^2(i^*)$, (b), of i^* to a set S of size j .

marginal contributions, so we start by defining them in terms of these marginal contributions and we later give their formal definition to show that they are well-defined. The marginal contributions of i^* are illustrated in Figure 2.

$$C_S^1(i) = \begin{cases} 1 & \text{if } |S| < L \\ 1 - \kappa & \text{otherwise} \end{cases}$$

$$C_S^2(i^*) = \begin{cases} 1 - \kappa & \text{if } |S| \leq H \\ 1 - \kappa - (|S| - H) \cdot \frac{1 - \kappa - (1 - \kappa)^2}{\sqrt{n}} & \text{if } H < |S| \leq H + \sqrt{n} \\ (1 - \kappa)^2 & \text{otherwise} \end{cases}$$

For $i \neq i^*$:

$$C_S^2(i) = \begin{cases} \frac{L - (1 - \kappa)}{L - 1} & \text{if } |S| < L - 1 \text{ or } (|S| = L - 1 \text{ and } i^* \in S) \\ 1 - \kappa & \text{if } (|S| = L - 1 \text{ and } i^* \notin S) \text{ or } L \leq |S| \leq H \text{ or} \\ & (H \leq |S| \leq H + \sqrt{n} \text{ and } i^* \notin S) \\ 1 - \kappa - \frac{1 - \kappa - (1 - \kappa)^2}{\sqrt{n}} & \text{otherwise} \end{cases}$$

The formal definitions of the functions are

$$C^1(S) = \begin{cases} |S| & \text{if } |S| < L \\ L + (|S| - L) \cdot (1 - \kappa) & \text{otherwise} \end{cases}$$

and

$$C^2(S) = \begin{cases} \frac{1_{i^* \in S} \cdot (1 - \kappa) + (|S| - 1_{i^* \in S}) \cdot \frac{L - (1 - \kappa)}{L - 1}}{L + (|S| - L) \cdot (1 - \kappa)} & \text{if } |S| < L \\ & \text{if } L \leq |S| \leq H \\ & \text{or } (H < |S| \leq H + \sqrt{n} \\ & \text{and } i^* \notin S) \\ L + (|S| - L) \cdot (1 - \kappa) \\ + 1 - \kappa - (|S| - H) \cdot \frac{1 - \kappa - (1 - \kappa)^2}{\sqrt{n}} & \text{if } H < |S| \leq H + \sqrt{n} \\ & \text{and } i^* \in S \\ L + (H + \sqrt{n} - L) \cdot (1 - \kappa) + 1_{i^* \in S} \cdot (1 - \kappa)^2 \\ + (|S| - 1_{i^* \in S} - (H + \sqrt{n})) \cdot (1 - \kappa - \frac{1 - \kappa - (1 - \kappa)^2}{\sqrt{n}}) & \text{otherwise} \end{cases}$$

The Shapley value of i^* with respect to C^1 and C^2 is then:

$$\phi_{i^*}^{C^1} = 1 \cdot \frac{1 - \epsilon'}{2} + (1 - \kappa) \cdot \frac{1 + \epsilon'}{2} \geq \frac{2 - \kappa}{2} - \epsilon$$

and

$$\phi_{i^*}^{C^2} \leq (1 - \kappa) \cdot \frac{1 + \epsilon'}{2} + (1 - \kappa)^2 \cdot \frac{1 - \epsilon'}{2} \leq \frac{(1 - \kappa)(2 - \kappa)}{2} + \epsilon$$

for an appropriate choice of ϵ' . Next, by Chernoff bound and a union bound, $L \leq |S| \leq H$ for polynomially many samples S from the uniform distribution, with probability $1 - e^{-\Omega(n)}$. Thus, $C^1(S) = C^2(S)$ for all samples S with probability $1 - e^{-\Omega(n)}$.

It remains to show that C^1 and C^2 are submodular with curvature κ , i.e., for any $S \subseteq T$ and $i \notin T$,

$$C_S(i) \geq C_T(i) \geq (1 - \kappa)C_S(i),$$

which is immediate for C^1 . Regarding C^2 , it is also immediate that $C_S^2(i^*) \geq C_T^2(i^*) \geq (1 - \kappa)C_S^2(i^*)$. For $i \neq i^*$, observe that

$$\frac{L - (1 - \kappa)}{L - 1} \leq 1 + \epsilon \quad \text{and} \quad 1 - \kappa - \frac{1 - \kappa - (1 - \kappa)^2}{\sqrt{n}} \geq 1 - \kappa - \epsilon,$$

so $C_S^2(i) \geq C_T^2(i) \geq (1 - \kappa - \epsilon)C_S^2(i)$. \square

Theorem 5. *There exists no constant $\alpha > 0$ such that coverage functions have Shapley value that is α -approximable from samples over the uniform distribution.*

Proof. Partition N into two parts A and B of equal size. Consider the following two functions:

$$C^1(S) = \begin{cases} 0 & \text{if } S = \emptyset \\ 1 & \text{if } |S \cap A| > 0, |S \cap B| = 0 \\ \frac{1}{\alpha^2} & \text{if } |S \cap A| = 0, |S \cap B| > 0 \\ 1 + \frac{1}{\alpha^2} & \text{otherwise} \end{cases}$$

$$C^2(S) = \begin{cases} 0 & \text{if } S = \emptyset \\ 1 & \text{if } |S \cap B| > 0, |S \cap A| = 0 \\ \frac{1}{\alpha^2} & \text{if } |S \cap B| = 0, |S \cap A| > 0 \\ 1 + \frac{1}{\alpha^2} & \text{otherwise} \end{cases}$$

These functions are coverage functions with $U = \{a, b_1, \dots, b_{1/\alpha^2}\}$ and $T_i = \{a\}$ or $T_i = \{b_1, \dots, b_{1/\alpha^2}\}$. By the Chernoff bound (Lemma 1) with $\delta = 1/2$ and $\mu = n/2$, if S is a sample from the uniform distribution, then

$$\Pr[|S \cap A| = 0] = \Pr[|S \cap B| = 0] < \Pr[|S \cap B| \leq n/4] \leq e^{-n/16},$$

so $C^1(S) \neq C^2(S)$ with probability at most $2e^{-n/16}$. It is easy to see that for any i , its Shapley value is either $2/n$ or $2/(\alpha^2 n)$ depending on which partition it is in. Combining this with Lemma 8 concludes the proof. \square

D Missing Analysis from Section 5

Theorem 6. *The data-dependent Shapley value is the unique value satisfying the four data-dependent axioms.*

We first show that if there exists a value satisfying the axioms, it must be the data-dependent Shapley value. Then, we show that the data-dependent Shapley value satisfies the axioms, which concludes the proof.

Lemma 9. *If there exists a value satisfying the four data-dependent Shapley axioms, then this value is the data-dependent Shapley value.*

Proof. Define \mathcal{D}_S to be the distribution such that $\Pr[S \sim \mathcal{D}_S] = 1$. Observe that the unique value satisfying the balance, symmetry, and zero element axioms must satisfy

$$\phi_i^{\mathcal{D}_S} = \begin{cases} \frac{C(S)}{|S|} & \text{if } i \in S \\ 0 & \text{otherwise.} \end{cases}$$

Since $\mathcal{D} = \sum_S \Pr[S \sim \mathcal{D}] \cdot \mathcal{D}_S$, the unique value satisfying the four axioms must satisfy $\phi_i^{\mathcal{D}} = \sum_S \Pr[S \sim \mathcal{D}] \cdot \phi_i^{\mathcal{D}_S} = \sum_{S:i \in S} \Pr[S \sim \mathcal{D}] \cdot C(S)/|S|$ where the first equality is by additivity and the second equality by the above observation. \square

Lemma 10. *The data-dependent Shapley value satisfies the four data-dependent Shapley axioms.*

Proof. We show that each axiom is satisfied.

- **Balance:** By definition, $\sum_{i \in N} \phi_i^{\mathcal{D}} = \sum_{i \in N} \sum_{S: i \in S} \Pr[S \sim \mathcal{D}] C(S)/|S|$, then by switching the order of the summations,

$$\begin{aligned} \sum_{i \in N} \sum_{S: i \in S} \Pr[S \sim \mathcal{D}] \frac{C(S)}{|S|} &= \sum_{S \subseteq N} \sum_{i \in S} \Pr[S \sim \mathcal{D}] \frac{C(S)}{|S|} \\ &= \sum_{S \subseteq N} \Pr[S \sim \mathcal{D}] C(S) \\ &= \mathbb{E}_{S \sim \mathcal{D}}[C(S)]. \end{aligned}$$

- **Symmetry:** Let $\mathcal{S}_i = \{S : i \in S, \Pr[S \sim \mathcal{D}] > 0\}$. If $\Pr_{S \sim \mathcal{D}}[|S \cap \{i, j\}| = 1] = 0$, then $\mathcal{S}_i = \mathcal{S}_j$ and

$$\phi_i^{\mathcal{D}} = \sum_{S \in \mathcal{S}_i} \Pr[S \sim \mathcal{D}] \frac{C(S)}{|S|} = \sum_{S \in \mathcal{S}_j} \Pr[S \sim \mathcal{D}] \frac{C(S)}{|S|} = \phi_j^{\mathcal{D}}.$$

- **Zero element:** If $\Pr_{S \sim \mathcal{D}}[i \in S] = 0$, then $\Pr[S \sim \mathcal{D}] = 0$ if $i \in S$. Thus, $\phi_i^{\mathcal{D}} = 0$.
- **Additivity:** By definition of ϕ and $\alpha\mathcal{D}_1 + \beta\mathcal{D}_2$,

$$\begin{aligned} \phi_i^{\alpha\mathcal{D}_1 + \beta\mathcal{D}_2} &= \sum_{S: i \in S} \Pr[S \sim \alpha\mathcal{D}_1 + \beta\mathcal{D}_2] \frac{C(S)}{|S|} \\ &= \alpha \sum_{S: i \in S} \Pr[S \sim \mathcal{D}_1] \frac{C(S)}{|S|} + \beta \sum_{S: i \in S} \Pr[S \sim \mathcal{D}_2] \frac{C(S)}{|S|} \\ &= \alpha\phi_i^{\mathcal{D}_1} + \beta\phi_i^{\mathcal{D}_2}. \end{aligned}$$

□

Theorem 7. *The empirical data-dependent Shapley value approximates the data-dependent Shapley value arbitrarily well, i.e.,*

$$|\tilde{\phi}_i^{\mathcal{D}} - \phi_i^{\mathcal{D}}| < \epsilon$$

with $\text{poly}(n, 1/\epsilon, 1/\delta)$ samples and with probability at least $1 - \delta$ for any $\delta, \epsilon > 0$.

Proof. Define $X_j = \begin{cases} \frac{C(S_j)}{|S_j|} & \text{if } i \in S_j \\ 0 & \text{otherwise} \end{cases}$ and observe that $(\sum_{j=1}^m X_j)/m = \tilde{\phi}_i^{\mathcal{D}}$ and

$\mathbb{E}[(\sum_{j=1}^m X_j)/m] = \phi_i^{\mathcal{D}}$. Clearly, $X_j \in [0, b]$ where $b := \max_S C(S)/|S|$, so by Hoeffding's inequality, $\Pr\left[|\tilde{\phi}_i^{\mathcal{D}} - \phi_i^{\mathcal{D}}| \geq \epsilon\right] \leq 2e^{-\frac{2m\epsilon^2}{\text{poly}(n)}}$ with $0 < \epsilon < 1$. □