Published online ahead of print November 14, 2012

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–25
ISSN 0364-765X (print) | ISSN 1526-5471 (online)

http://dx.doi.org/10.1287/moor.1120.0567

Optimal Cost Sharing for Resource Selection Games

Philipp von Falkenhausen
Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, Germany, falkenhausen@math.tu-berlin.de

Tobias Harks
School of Business and Economics, Maastricht University, 6211 Maastricht, The Netherlands,
t.harks@maastrichtuniversity.nl

Joint use of resources with usage-dependent cost raises the question: who pays how much? We study cost sharing in resource selection games where the strategy spaces are either singletons or bases of a matroid defined on the ground set of resources. Our goal is to design cost sharing protocols so as to minimize the resulting price of anarchy and price of stability. We investigate three classes of protocols: basic protocols guarantee the existence of at least one pure Nash equilibrium; separable protocols additionally require that the resulting cost shares only depend on the set of players on a resource; uniform protocols are separable and require that the cost shares on a resource may not depend on the instance, that is, they remain the same even if new resources are added to or removed from the instance. We find optimal basic and separable protocols that guarantee the price of stability and price of anarchy to grow logarithmically in the number of players, except for the case of matroid games induced by separable protocols where the price of anarchy grows linearly with the number of players. For uniform protocols we show that the price of anarchy is unbounded even for singleton games.

Key words: resource selection games; cost sharing; matroids; congestion games; cost functions
MSC2000 subject classification: Primary: 91A10; secondary: 91A43, 90B10, 68W01
OR/MS subject classification: Primary: games, noncooperative; secondary: mathematics, functions
History: Received January 26, 2012; revised August 1, 2012. Published online in *Articles in Advance*.

1. Introduction. We study resource selection games, where a set of resources is given and the space of pure strategies of a player consists of a set of subsets of the resources. The cost of a resource is a nondecreasing function of the sum of the demands of players choosing the resource. Classical examples of resource sharing games are Rosenthal's congestion games (Rosenthal [50]), where the cost of a resource depends only on the number of players choosing the resource, and the private cost of a player is the sum of the *average costs* of the chosen resources.

In contrast to congestion games, we focus in this paper on applications where the load dependent cost of a resource is money that can be shared arbitrarily among the players. One such application arises in network design games modeling the interaction of selfish players jointly designing a network infrastructure (see Anshelevich et al. [6], Chen et al. [18]). In a network design game, the resources correspond to edges in a directed or undirected graph and each player wants to establish a path or a spanning tree satisfying a player-specific bandwidth requirement. Although Chen et al. [18] assumed that edges have fixed costs with unbounded capacity (or nondecreasing concave costs as in Anshelevich et al. [6]), a more realistic model for satisfying multiple bandwidth requirements on an edge is the so-called *cable* model, where integer multiples of a cable with a certain capacity can be bought at a fixed price per increment (Antonakopoulos et al. [7]). More complicated types of cost functions may arise when there are different cable types that have a fixed cost per increment plus a load dependent cost function that only holds for a particular capacity interval. Given the resource cost functions and bandwidth requirements, in an ideal solution resources are allocated so as to minimize the aggregated resource costs. In decentralized systems, however, players will selfishly select resources/edges for their demands based on the cost shares they have to pay. Hence, the rules by which the cost of a resource is shared among its users play a key role as they determine the equilibrium states of the strategic game induced.

Our goal in this paper is to define *cost sharing protocols* such that pure Nash equilibria (PNE for short) of the induced strategic games always exist and the efficiency loss caused by selfish resource selection is minimized. Koutsoupias and Papadimitriou [41] and later Anshelevich et al. [6] introduced measures to quantify this efficiency loss known as the *price of anarchy* and the *price of stability*. The price of anarchy (PoA) is defined as the worst-case ratio of the cost of a Nash equilibrium over the cost of a system optimum, while the price of stability (PoS) captures the ratio of the best possible Nash equilibrium over a system optimum. We focus on two combinatorial structures of strategy spaces. In the first case, we study singleton resource selection games (singleton games for short), where every player uses exactly one resource. This class of games has applications in scheduling where each player is associated with a job of nonnegative weight. The job can be processed on every resource, and the monetary cost on a resource (for instance energy costs as in Yao et al. [57]) is a nondecreasing function of its total load. In the second case, we lift the assumption of singleton strategies to the more general setting of bases of matroids (matroid games for short). The strategy space of every player is the set of bases of some matroid defined on the ground set of the resources. Matroids have a rich combinatorial structure and include, for instance, the aforementioned class of network design games, where each player wants to allocate a spanning tree in a graph.

For both scenarios, we assume that the private cost of a player is determined by a cost sharing protocol, i.e., the protocol determines how the cost of each resource is shared among its users. For instance, a simple protocol that has been analyzed in the literature on congestion games is *average* or *proportional cost sharing* (see Aland et al. [3], Awerbuch et al. [9], Bhawalkar et al. [10], Correa et al. [22], and Roughgarden and Tardos [52]). Before we can answer the question which protocol induces games with the best equilibrium states, we first have to precisely define the design space of *feasible* cost sharing protocols. We impose the following four design assumptions that are defined more formally in §2. These properties have been introduced first by Chen et al. [18] in the context of the design of cost sharing protocols for fixed-cost network design games (a formal definition of these games will be given in §1.3).

(i) *Budget balance*. For every outcome of a game induced by the cost sharing protocol, the cost of each resource is exactly covered by the collected cost shares of the players using the resource.

(ii) *Stability*. There is at least one pure strategy Nash equilibrium in each game induced by the cost sharing protocol.

(iii) *Separability*. When assigning the cost shares on a given resource, the protocol has no information about the load on other resources.

(iv) *Uniformity*. When assigning the cost shares on a given resource, the protocol has no information about the existence of other resources.

A cost sharing protocol is called *basic* if it satisfies (i)–(ii), *separable* if it satisfies (i)–(iii), and *uniform* if it satisfies (i)–(iv). We briefly discuss the four properties and refer to Chen et al. [18] for a more detailed treatment. The condition (i) is the least controversial in the context of cost sharing protocols. The stability condition (ii) requires the existence of at least one Nash equilibrium in pure strategies. Although this requirement restricts the search space for cost sharing protocols, it is certainly the solution concept of choice when mixed or correlated strategies have no meaningful physical interpretation in the game played; see also the discussion in Osborne and Rubinstein [48, §3.2] about critics of mixed Nash equilibria. Although condition (iii) seems restrictive, it is crucial for practical applications in which cost sharing protocols have only local information about their own resource usage (see for instance the TCP/IP protocol design, where routers drop packets based on some function of the number of packets in the queue (Srikant [55])). Uniformity (iv) is the strongest and perhaps the most problematic design restriction. A uniform protocol is not only separable but also *strongly local* in the sense that the cost shares of a resource are independent of the set of resources available to the game designer. This property may be crucial for systems in which the resources can be added or removed over time and a reconfiguration of the system (changing the cost sharing protocol) is too costly.

1.1. Our results. We systematically analyze the achievable price of anarchy and stability by basic, separable, and uniform cost sharing protocols in the context of singleton games and matroid games. Whereas the price of anarchy and stability constitutes a worst-case measure for the inefficiency of pure Nash equilibria across all instances, we also address the problem of designing *universally* optimal protocols that are optimal for every instance.

1.1.1. Results for singleton games. For singleton games, we prove that among all basic and separable protocols, there is an optimal protocol minimizing the resulting price of anarchy and price of stability *simultaneously*. For n-player singleton games, the optimal value of the price of anarchy and stability is precisely the nth harmonic number $\mathscr{H}_n = \sum_{i=1}^{n}(1/i)$. As a key element of the proof, we obtain a complete characterization of pure Nash equilibria that can be induced by a basic or separable protocol. We also derive sufficient conditions for a strategy profile to be the most expensive pure Nash equilibrium. Our proof of this result is constructive by providing a cost sharing protocol that induces a strategy profile as the most expensive pure Nash equilibrium provided that it satisfies these conditions. We then show that this protocol gives rise to an optimal cost sharing protocol simultaneously minimizing the price of anarchy and stability as mentioned above. Our characterization of pure Nash equilibria can further be used to design a universally optimal separable protocol minimizing the cost of an achievable pure Nash equilibrium for *every* instance. For uniform cost sharing protocols we show that they cannot guarantee a bounded price of anarchy. We construct a lower bound involving a family of instances with only three players, at most three resources, and cost functions with nondecreasing costs per unit.

1.1.2. Results for matroid games. Matroid games are a generalization of singleton games and consequently all lower bounds obtained for singleton games carry over to the matroid setting. For separable and basic protocols,

TABLE 1. Overview about the price of stability and price of anarchy results.

	Singleton games		Matroid games	
	PoS	PoA	PoS	PoA
Uniform	$\geq\mathscr{H}_n$	∞	$\geq\mathscr{H}_n$	∞
Separable	$\mathscr{H}_n$	$\mathscr{H}_n$	$\mathscr{H}_n$	n
Basic	$\mathscr{H}_n$	$\mathscr{H}_n$	$\mathscr{H}_n$	$\mathscr{H}_n$

we devise an optimal (separable) protocol minimizing the price of stability resulting in a tight worst-case bound of $\mathscr{H}_n$. We again prove this result by obtaining a complete characterization of pure Nash equilibria that can be induced by a separable protocol. This characterization crucially relies on the unique resource exchange properties of matroids. We can again use this characterization to obtain a universally optimal separable protocol minimizing the cost of an achievable pure Nash equilibrium for every instance. In contrast to singleton games, however, our characterization does not carry over to basic protocols. In fact, we find a structural difference of the achievable price of anarchy when going from separable to basic protocols: we devise an optimal separable protocol with price of anarchy of exactly n. For the larger class of basic protocols we devise an optimal basic protocol having price of anarchy of $\mathscr{H}_n$. All our results are summarized in Table 1.

1.2. Significance and techniques used. Our work is closely related to the paper by Chen et al. [18]. In their paper, the authors study the design of cost sharing protocols for fixed-cost network design games. In a network design game, each player i wishes to send a (unit) demand along a path in a (directed or undirected) network, connecting her source node s_i to her terminal node t_i. Every edge has a fixed cost and the goal is to design a separable or uniform cost sharing protocol so as to minimize the resulting price of anarchy and stability. Our approach follows their lead in terms of the protocol design perspective and the feasible protocol space, but we apply cost sharing protocols to the structurally different class of singleton and matroid models. Matroid games include the class of network design games, where instead of single s-t paths, players wish to send their demand along a spanning tree. In contrast to previous works on network design games, our model allows for the first time arbitrary nondecreasing cost functions instead of fixed costs (or concave costs) on the resources. This way we are able to model more realistic cost structures occurring in network design. Typical cost functions are step functions (see Figure 1 (left)), where every cost level corresponds to a different cable type that can be installed (cf. Antonakopoulos et al. [7]). Andrews et al. [5] recently introduced a network design problem in telecommunications, where cost functions with "diseconomies of scale" are used to model the cost accounting for energy consumption when routers apply speed scaling to process packets. The proposed cost function is defined as $c_r(\ell) = \sigma + \delta \cdot \ell^{\alpha}$, $\alpha > 1$, σ, $\delta > 0$ if $\ell > 0$ and, $c_r(0) = 0$; see Figure 1 (right) for an illustration. Clearly, this function and also the previous function are neither concave nor convex.

A central challenge in cost sharing protocol design involving arbitrary nondecreasing cost functions is to ensure that the induced strategic games always posses a PNE. Because prevailing approaches for proving the existence of PNE (such as potential functions and fixed-point theorems) are not directly applicable to our games we develop a new approach to cope with the equilibrium existence problem: we exploit structural properties of PNE in resource selection games to derive a complete (protocol independent) characterization of strategy profiles that can be obtained as a PNE. We call such a strategy profile *decharged*. Informally, a strategy profile is

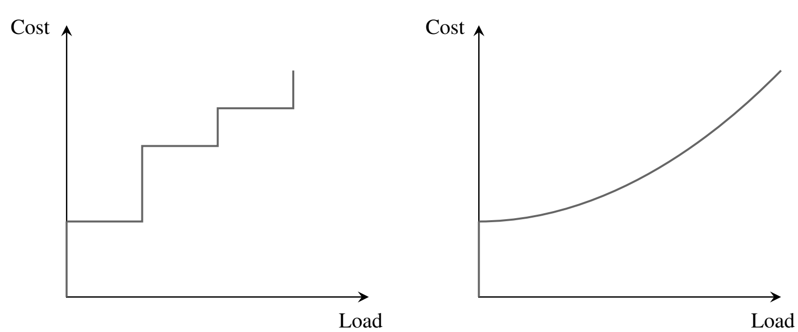

FIGURE 1. Cost functions with nonconcave/nonconvex cost functions.

decharged if the cost on every resource is less than the sum of the costs that arise if every player on the resource exchanges this resource with her best alternative. We then devise a family of protocols (x-enforcing protocols) parameterized by a decharged profile x. Using the notion of decharged profiles, we devise an algorithm that takes the optimal strategy as input and iteratively modifies the optimal profile until it is decharged. For each iteration of the algorithm we keep track of the resulting cost of the intermediate strategy profiles and prove that the final strategy profile is at most a factor $\mathscr{H}_n$ away from the optimal cost. Thus, we use our characterization not only to show the existence of PNE but also to determine the cost of PNE and hence the price of stability. This proof technique is quite different from the prevailing approaches of combining the Nash inequality with the potential function method (cf. Christodoulou and Koutsoupias [19], Correa et al. [22]), or by arguing directly that the optimal solution can be a pure Nash equilibrium (cf. Anshelevich et al. [6], Epstein et al. [23], Harks et al. [32]).

For obtaining our bounds on the price of anarchy, we develop a new approach that is different from the prevailing "smoothness" arguments (cf. Aland et al. [3], Bhawalkar et al. [10], Harks [30], Roughgarden [51]) relating the cost of a worst-case Nash equilibrium via the Nash inequalities to a weighted average of the cost of an optimal strategy and the cost of the worst-case Nash equilibrium itself. Instead, in our approach, we first relate the cost of a worst-case Nash equilibrium to the payments of the players if they individually deviated to the strategy of the decharged profile returned by the algorithm. We then use properties of the protocol, the decharged condition, and the algorithm to estimate the resulting term against the cost of an optimal profile achieving a tight bound. Although this part of the paper is perhaps the most involved, we are confident that our approach can be useful if standard approaches for bounding the price of anarchy fail.

1.3. Further related work. In the previous section we described structural differences (weighted players, arbitrary nondecreasing cost functions, and symmetric strategy spaces) between our work and the paper by Chen et al. [18]. These structural differences also apply to other works on network design games (cf. Anshelevich et al. [6], Bilò et al. [11], Chekuri et al. [15], Chen and Roughgarden [16]) and result in different approaches and different achievable bounds. For example, whereas Chen et al. [18] proved bounds on the price of anarchy for uniform protocols of order $\Theta(\log(n))$, $\Theta(\text{polylog}(n))$, and n for undirected single-sink instances, undirected multicommodity instances, and directed single-sink instances, respectively, we show that in our model even for singleton games such results are impossible. The price of anarchy for uniform protocols inducing singleton games is unbounded. Finally, it is worth noting that, whereas Chen et al. [18] analyzed separable and uniform protocols, we additionally analyze the larger class of basic protocols.

There is a large body of work on scheduling games (or singleton congestion games) with unweighted and weighted players (Ackermann et al. [2], Even-Dar et al. [24], Fotakis et al. [27], Gairing et al. [28], Ieong et al. [36], Milchtaich [42]). Most of these papers study the existence and price of anarchy of pure Nash equilibria for the *proportional* cost sharing protocol in which the private cost of every player is equal to the cost on the resource, or equivalently, the cost share of every player is proportional to its own demand, that is, if player i has a demand of $d_i > 0$, then she pays $d_r(\ell_r) \cdot d_i$, where d_r is a nondecreasing delay function. These works, however, do not consider the design perspective of cost sharing protocols. For weighted congestion games with arbitrary strategy spaces, proportional cost sharing has the severe drawback that PNEs need not exist (unless the functions $d_r(\ell_r)$ are either affine or exponential (Fotakis et al. [25, 26], Goemans et al. [29], Harks and Klimm [31], Milchtaich [43], Panagopoulou and Spirakis [49]), or the strategy spaces have a specific combinatorial structure, e.g., matroid games (Ackermann et al. [2])). To address this drawback, Kollias and Roughgarden [40] proposed a cost sharing protocol based on the Shapley value for which they are able to prove existence of PNE. They further prove price of anarchy/stability results for polynomial cost per unit functions with nonnegative coefficients. It is unknown, however, whether this protocol is optimal even for the special class of polynomial cost per unit functions. For further work on congestion games with proportional sharing assuming nonincreasing marginal cost functions (modeling economies of scale or buy at bulk) we refer to Albers [4], Anshelevich et al. [6], Epstein et al. [23], Hoefer [35], and Rozenfeld and Tennenholtz [53].

Christodoulou et al. [20] and follow-up papers (such as Caragiannis [13], Cole et al. [21], Immorlica et al. [37]) study *coordination mechanisms* and their price of anarchy in scheduling games in which n players assign a task to one of m machines. Rather than paying a share of the resulting cost of a machine as in our scenario, the players in these games consider the completion time of their respective job as private cost. Hence, although they assume a mechanism/protocol design perspective as we do, the class of games analyzed is quite different.

Sharing the cost of resources from the perspective of cooperative game theory is a central topic in economics and operations research (cf. Archer et al. [8], Bogomolnaia et al. [12], Moulin and Shenker [47], or the survey of Moulin [44] with a pointer to further references). There is also a large body of papers studying cost sharing from a noncooperative perspective (cf. Chen and Zhang [17], Harks and Miller [33], Johari and Tsitsiklis [38, 39], Moulin [45, 46]). The model considered in these papers is different to ours as they all assume a compact and convex strategy space together with convexity assumptions on the feasible cost sharing functions.

1.4. Structure of the paper. In §2 we present the model together with our key assumptions on feasible cost sharing protocols used in this paper. We first start in §3 with the more intuitive setting of singleton games on parallel resources. In §4, we then turn to the more general setting of matroids. We conclude the paper in §5 by pointing out open questions and new research directions.

2. Model and problem statement. A *resource selection model* is represented by a tuple (N, M, Σ, d, c), where $N = \{1, \ldots, n\}$ is the set of players and $M = \{a_1, \ldots, a_m\}$ the set of resources. Each player i chooses a subset X_i of these resources contained in her strategy set $\Sigma_i \subseteq 2^M$. We denote by $X := (X_1, \ldots, X_n)$ the strategy profile across all players and correspondingly $\Sigma := \prod_{i\in N} \Sigma_i$. Unless stated otherwise, we assume that $\Sigma_i = \Sigma_j$ for all $i, j \in N$, that is, the strategy spaces are symmetric. The vector $d = (d_i)_{i\in N}$ specifies the player's weights, which in a strategy profile X sum up on each resource $a \in M$ to the load $\ell_a(X) := \sum_{i\in S_a(X)} d_i$, where $S_a(X) := \{i\colon a \in X_i\}$ is the set of players using a. The resources' cost functions are given by the vector $c = (c_a)_{a\in M}$, they are nondecreasing in the load. We define the social cost of profile X as $C(X) := \sum_{a\in M} c_a(\ell_a(X))$. Abusing notation, we often refer to the cost on resource a by $c_a(X)$. We now give two concrete scenarios of resource selection models that we study in this paper.

Example 2.1 (Singleton Games). In a singleton game (or scheduling game), every player i chooses exactly one resource, that is $|X_i| = 1$ for all $X_i \in \Sigma_i$ and $i \in N$. A natural interpretation is that each player has a job of weight d_i and chooses a resource to process its job.

We now give a more complex scenario that captures the situation where players want to build a spanning tree in an undirected graph.

Example 2.2 (MST Games). We are given an undirected graph $G = (V, E)$ with nonnegative and nondecreasing edge cost functions $c_e(\ell)$, $e \in E$. In a minimum spanning tree (MST) game, every player i is associated with demand of size $d_i > 0$ and routes its demand along a spanning tree. Formally, we set $M = E$ and the sets X_i, $i \in N$, are the spanning trees of G.

We study how different ways of sharing the costs of a resource affect the resulting pure Nash equilibria of the induced game. To model this, we introduce cost sharing protocols Ξ that assign cost share functions $\xi_{i,a}\colon \Sigma \to \mathbb{R}$ for all $i \in N$ and $a \in M$ to the resource selection model (N, M, Σ, d, c) and thus induce the strategic game (N, Σ, ξ). For a player i, her total private cost is $\xi_i(X) := \sum_{a\in X_i} \xi_{i,a}(X)$ and we assume that every player strives to minimize her private cost. An important solution concept in noncooperative game theory are pure Nash equilibria. Using standard notation in game theory, for a strategy profile $X \in \Sigma$ we denote by

$$(Z_i, X_{-i}) := (X_1, \ldots, X_{i-1}, Z_i, X_{i+1}, \ldots, X_n) \in \Sigma$$

the profile that arises if only player i deviates to strategy $Z_i \in \Sigma_i$.

Definition 2.1 (Pure Nash Equilibrium). Let (N, Σ, ξ) be a strategic game. The profile X is a pure Nash equilibrium if no player i can strictly reduce her private cost by unilaterally moving to a different strategy, that is, for all $i \in N$

$$\xi_i(X) \le \xi_i(Z_i, X_{-i}) \quad \text{for all } Z_i \in \Sigma_i.$$

Two well-established concepts that quantify the efficiency of Nash equilibria are the *price of anarchy* and the *price of stability*. The price of anarchy measures the largest possible ratio of the cost of a Nash equilibrium and the cost of an optimal profile. The price of stability measures the smallest ratio of the cost of a Nash equilibrium and the cost of an optimal profile. For a cost sharing protocol Ξ, we define by $\mathrm{PoA}(\Xi)$ and $\mathrm{PoS}(\Xi)$ the corresponding worst-case price of anarchy and price of stability across games induced by protocol Ξ. The main goal of this paper is to design cost sharing protocols that minimize the price of anarchy and price of stability, respectively. Of course, the attainable objective values crucially depend on the design space that we permit. The following properties have been first proposed by Chen et al. [18] in the context of designing cost sharing protocols for network design games.

Definition 2.2 (Properties of Cost Sharing Protocols). A cost sharing protocol Ξ is

(i) *stable* if it induces only games that admit at least one pure Nash equilibrium;

(ii) *basic* if it is stable and additionally *budget balanced*, i.e., if it assigns all resource selection models (N, M, Σ, d, c) with cost share functions $\xi_{i,a}$ such that for all $a \in M$ and $X \in \Sigma$

$$c_a(X) = \sum_{i\in S_a(X)} \xi_{i,a}(X) \qquad \text{and} \qquad \xi_{i,a}(X) = 0 \quad \text{for all } i \notin S_a(X)$$

(this property requires $c_a(0) = 0$ for unused resources, which we will assume in the paper);

(iii) *separable* if it is basic and if it induces only games for which in any two profiles X, $X' \in \Sigma$ for every resource $a \in M$,

$$S_a(X) = S_a(X') \Rightarrow \xi_{i,a}(X) = \xi_{i,a}(X') \quad \text{for all } i \in S_a(X);$$

(iv) *uniform* if it is separable and if it assigns any two models (N, M, Σ, d, c), (N, M', Σ', d, c') with cost share functions $\xi_{i,a}$ and $\xi'_{i,a}$ such that the following condition holds. For all $a \in M \cap M'$ with $c_a = c'_a$ and all profiles $X \in \Sigma$, $X' \in \Sigma'$

$$S_a(X) = S_a(X') \Rightarrow \xi_{i,a}(X) = \xi'_{i,a}(X') \quad \text{for all } i \in S_a(x).$$

Informally, separability means that in a profile X the values $\xi_{i,a}(X)$, $i \in N$ depend only on the set $S_a(X)$ of players sharing resource a and disregard all other information contained in X. Still, separable protocols can assign cost share functions that are specifically tailored to the given resource selection model, for example based on an optimal profile. Uniform protocols are not allowed to do this, they even disregard the layout of the model and assign the same cost shares when resources are added to or removed from the model.

We denote by $\mathcal{B}_n$, $\mathcal{S}_n$, and $\mathcal{U}_n$ the set of basic, separable, and uniform protocols for resource selection games with n players, respectively. We obtain the following optimization problems that we address in this paper:

$$\min_{\Xi \in \mathcal{B}_n} \text{PoA}(\Xi), \quad \min_{\Xi \in \mathcal{B}_n} \text{PoS}(\Xi), \quad \min_{\Xi \in \mathcal{S}_n} \text{PoA}(\Xi), \quad \min_{\Xi \in \mathcal{S}_n} \text{PoS}(\Xi),$$
$$\min_{\Xi \in \mathcal{U}_n} \text{PoA}(\Xi), \quad \text{and} \quad \min_{\Xi \in \mathcal{U}_n} \text{PoS}(\Xi).$$

3. Singleton games. This section deals with games such as the one in Example 2.1, where players choose a single resource instead of multiple resources. More precisely, we call a resource selection model a singleton model if for every player i the strategy set is exactly the set of available resources, i.e., $\Sigma_i = M$. This allows us to simplify notation for this section: we denote a singleton model by (N, M, d, c) and we denote a player's strategy by a single resource x_i (instead of a set of resources X_i). Consequently, we denote profiles as vectors of n resources $x \in M^n$. Budget balanced protocols have $\xi_{i,a}(x) = 0$ for all $a \neq x_i$ and, hence, $\xi_i(x) = \xi_{i,x_i}(x)$, which allows us to define these protocols by aggregate cost share functions ξ_i instead of the per resource cost share functions $\xi_{i,a}$.

In the following, we start with studying basic and separable protocols. We find structural properties of Nash equilibria of the games induced by such protocols that let us construct an optimal protocol. Throughout this section, the players are assumed to be ordered by nondecreasing weights:

$$d_1 \leq d_2 \leq \cdots \leq d_n. \tag{1}$$

3.1. Characterization of Nash equilibria for basic and separable protocols. To design protocols that have a cheap Nash equilibrium, it is crucial to understand structural properties of Nash equilibria. We find a complete characterization of Nash equilibria in singleton games induced by basic and separable protocols. First, let us introduce some notation.

Definition 3.1 (Weakly Decharged Profile). Consider a singleton model (N, M, d, c). A resource $a \in M$ is *weakly decharged* in a profile $x \in M^n$ if

$$c_a(x) \leq \sum_{i \in S_a(x)} \min_{b \in M} c_b(b, x_{-i}).$$

The profile x itself is called *weakly decharged* if all resources are weakly decharged.

We further introduce the x-enforcing protocol.

Definition 3.2 (x-Enforcing Protocol). The x-*enforcing protocol* takes as input a weakly decharged profile x. We use x to define for any profile z and resource a the set of *foreign players on a* $S^1_a(z) := \{i \in S_a(z) \setminus S_a(x)\}$

and the set of *strong foreign players on a* $S_a^2(z) := \{i \in S_a(z) \backslash S_a(x)\colon c_{x_i}(x) = 0\}$. Then, the x-enforcing protocol assigns for all $i \in N$, $z \in M^n$ the cost share functions

$$\xi_i(z) := \begin{cases} \dfrac{\min_{b\in M} c_b(b, x_{-i})}{\sum_{j\in S_{x_i}(x)} \min_{b\in M} c_b(b, x_{-j})} \cdot c_{x_i}(x), & \text{if } S_{z_i}(z) = S_{z_i}(x) \text{ and } c_{x_i}(x) > 0,\\ c_{z_i}(z), & \text{if } S_{z_i}^2(z) \neq \varnothing \text{ and } i = \min S_{z_i}^2(z),\\ c_{z_i}(z), & \text{if } S_{z_i}^2(z) = \varnothing, S_{z_i}^1(z) \neq \varnothing \text{ and } i = \min S_{z_i}^1(z),\\ c_{z_i}(z), & \text{if } S_{z_i}^1(z) = \varnothing, S_{z_i}(z) \subset S_{z_i}(x) \text{ and } i = \min S_{z_i}(z),\\ 0, & \text{else.} \end{cases}$$

Informally, if $S_a(z) = S_a(x)$, the players on resource a share the cost proportional to their opportunity cost (cost of change) in profile x. Otherwise, the smallest strong foreign player (deviating from x although $\xi_{i,a}(x) = 0$), foreign player (deviating from x), or home player (not deviating from x) pays the entire cost of the resource. Observe that in weakly decharged profiles x we have

$$\sum_{j\in S_a(x)} \min_{b\in M} c_b(b, x_{-j}) > 0 \quad \text{for all } a \in M \text{ with } c_a(x) > 0$$

and thus the protocol is well defined. We are now ready to state our first characterization.

Theorem 3.1. *For any singleton model* (N, M, d, c) *and profile* x, *the following statements are equivalent.*
(i) *the profile* x *is weakly decharged,*
(ii) *the profile* x *is a pure Nash equilibrium in the game induced by some basic protocol, and*
(iii) *the profile* x *is a pure Nash equilibrium in the game induced by some separable protocol.*

Observe that (iii) ⇒ (ii) holds because by definition separable protocols are a subclass of basic protocols. We prove (i) ⇒ (iii) and (ii) ⇒ (i) by two lemmas.

Lemma 3.1. *For every weakly decharged profile* x, *the* x*-enforcing protocol is a separable protocol and* x *is a pure Nash equilibrium in the induced game.*

Proof. Budget balance and separability follow immediately from the definition of the protocol. Thus, we prove only that x is a pure Nash equilibrium in the induced game. For all resources $a \in M$ with $c_a(x) > 0$ we are in the first case of the definition of the protocol, thus, we obtain for all $i \in S_a(x)$

$$\xi_i(x) = \frac{\min_{b\in M} c_b(b, x_{-i})}{\sum_{j\in S_a(x)} \min_{b\in M} c_b(b, x_{-j})} \cdot c_a(x) \leq \min_{b\in M} c_b(b, x_{-i}) \leq \min_{b\in M\backslash\{a\}} \xi_i(b, x_{-i}),$$

where the first inequality holds because profile x is weakly decharged. For all other resources $a \in M$, we have $\xi_i(x) = c_a(x) = 0$ for all $i \in S_a(x)$ and thus x is a pure Nash equilibrium. □

We prove (ii) ⇒ (i) from Theorem 3.1 by the following lemma.

Lemma 3.2. *Consider the game induced by some basic protocol on some singleton model* (N, M, d, c). *Then, any pure Nash equilibrium* x *is weakly decharged.*

Proof. If x is a pure Nash equilibrium, then $\xi_i(x) \leq \min_{b\in M} \xi_i(b, x_{-i})$ for all $i \in N$ and, hence, because of budget balance of the protocol,

$$c_a(x) = \sum_{i\in S_a(x)} \xi_i(x) \leq \sum_{i\in S_a(x)} \min_{b\in M} \xi_i(b, x_{-i}) \leq \sum_{i\in S_a(x)} \min_{b\in M} c_b(b, x_{-i})$$

for all resources $a \in M$. Thus, x is weakly decharged. □

The above characterization reduces the problem of finding the basic/separable protocol with lowest PoS to the minimization problem

$$\min_{x\in M^n} C(x) \qquad \text{s.t.} \quad c_a(x) \leq \sum_{i\in S_a(x)} \min_{b\in M} c_b(b, x_{-i}) \quad \forall\, a \in M,$$

that is, to finding the cheapest weakly decharged profile. As the characterization is independent of our social welfare function $C(x)$, similar minimization problems can be formulated when designing PoS-optimal basic/separable protocols for arbitrary objective functions.

Corollary 3.1. *Let (N, M, d, c) be a singleton model and let $F\colon M^n :\to \mathbb{R}$ be a social welfare function. Then,* $\min_{\Xi\in\mathcal{B}_n}\mathrm{PoS}(\Xi; F)$ *and* $\min_{\Xi\in\mathcal{S}_n}\mathrm{PoS}(\Xi; F)$ *can be reduced to solving the optimization problem*

$$\min_{x\in M^n} F(x) \qquad \text{s.t.} \quad c_a(x) \le \sum_{i\in S_a(x)} \min_{b\in M} c_b(b, x_{-i}) \quad \forall\, a\in M.$$

The x-enforcing protocol is, given the cheapest decharged profile x, universally optimal for any social welfare function F. In other words, it is not only optimal from a worst-case perspective, but even on a per instance perspective.

Before we show how to find cheap decharged profiles and estimate their cost in §3.2.1, we first establish a corresponding structural result for the PoA.

3.2. Criterion for worst-case Nash equilibria for basic and separable protocols. The characterization of pure Nash equilibria will allow us to design a PoS-optimal protocol, but for a PoA-optimal protocol we need to go a step further and ask which profiles are the most expensive pure Nash equilibria of games induced by basic and separable protocols. We derive a sufficient condition for this, using the x-enforcing protocol from the previous section.

Definition 3.3 (Strongly Decharged Profile). Consider a singleton model (N, M, d, c). A resource $a\in M$ is *strongly decharged* if it is weakly decharged and additionally

$$c_a(x) < \sum_{i\in S_a(x)} \min_{b\in M} c_b(b, x_{-i}), \quad \text{if } |S_a(x)| > 1 \text{ and } c_a(x) > 0. \tag{2}$$

Resources that are not strongly decharged are called *charged*. The profile x is called *strongly decharged* if all resources are strongly decharged.

Theorem 3.2. *Let (N, M, d, c) be a singleton model and x be a strongly decharged profile. Then x is the most expensive pure Nash equilibrium in the game induced by the x-enforcing protocol.*

Proof. We show that for any pure Nash equilibrium $z\ne x$ we have $C(z)\le C(x)$. To this end, fix such a z and let $i := \min\{j\in N\colon z_j\ne x_j\}$ be the smallest player who deviates from x. First, note that for all $j>i$,

$$\xi_j(z) \le \begin{cases} \xi_j(z_i, z_{-j}) = 0, & \text{if } c_{x_j}(x) > 0\\ \xi_j(x_j, z_{-j}) = 0, & \text{if } c_{x_j}(x) = 0,\end{cases} \tag{3}$$

because z is a pure Nash equilibrium. Hence,

$$c_a(z) = 0 \quad \text{for all resources } a\ne z_i \text{ with foreign players } S_a^1(z)\ne\varnothing. \tag{4}$$

Also, $c_a(z)\le c_a(x)$ for all resources $a\ne z_i$ that only have home players $S_a^0(z) = S_a(z)$, because for these resources $\ell_a(z)\le \ell_a(x)$. Thus, we already have

$$c_a(z)\le c_a(x) \quad \text{for all resources } a\ne z_i. \tag{5}$$

If there is a strong foreign player on z_i, then even $c_{z_i}(z)=0$ and we are done. Thus, from now on we assume that there are no strong foreign players on z_i. We can bound $c_{z_i}(z)$ from above using the Nash inequality $c_{z_i}(z)=\xi_i(z)\le \xi_i(x_i, z_{-i})$. The remaining proof focuses on bounding the value $\xi_i(x_i, z_{-i})$ from above.

The value of $\xi_i(x_i, z_{-i})$ assigned by the x-enforcing protocol depends on $S_{x_i}(x_i, z_{-i})$ and $c_{x_i}(x)$, for which there are three possibilities, corresponding to the first, fourth, and fifth case from the definition of the x-enforcing protocol. These cases are

(i) $S_{x_i}(x_i, z_{-i}) = S_{x_i}(x)$ and $c_{x_i}(x)>0$, where the protocol returns $\xi_i(x_i, z_{-i}) = \xi_i(x)$;
(ii) $S_{x_i}(x_i, z_{-i})\subset S_{x_i}(x)$ and $i=\min S_{x_i}(x_i, z_{-i})$, where the protocol returns $\xi_i(x_i, z_{-i}) = c_{x_i}(x_i, z_{-i})$; and
(iii) all cases in which the protocol returns $\xi_i(x_i, z_{-i})=0$.

In each case we will find $c_{x_i}(z)+c_{z_i}(z)\le c_{x_i}(x)+c_{z_i}(x)$ and thus with (5) we have $C(z)\le C(x)$, which proves the Theorem. Note that (5) already implies $c_{x_i}(z)\le c_{x_i}(x)$.

We begin with Case (i). The condition $c_{x_i}(x) > 0$ implies that if there is some strong foreign player $j > i$ (with $z_j \neq x_j$ and $c_{x_j}(x) = 0$), then $c_{z_i}(z) = \xi_i(z) \leq \xi_i(z_j, z_{-i}) = 0$ and we are done. Thus, we will in the following assume that there are no strong foreign players at all. If $\xi_i(x) = 0$, we obtain $0 = \xi_i(x) = \xi_i(x_i, z_{-i}) \geq \xi_i(z) = c_{z_i}(z)$, because we are in Case (i). Thus, we will also assume

$$\xi_i(x) > 0. \tag{6}$$

We now compare the allocation of load in the profiles z and x, respectively. First, we consider resources $a \neq z_i$, which host foreign players $j \in S_a(z) \backslash S_a(x)$. For these foreign players we obtain

$$\begin{aligned} \min_{b\in M} c_b(b, x_{-j}) &\geq \min_{b\in M} c_b(b, x_{-i}) && \text{(7a)} \\ &\geq \xi_i(x) && \text{(7b)} \\ &> 0. && \text{(7c)} \end{aligned}$$

We chose player i such that $j > i$ for all other foreign players j and hence (by (1)) $d_j \geq d_i$. As the cost functions are nondecreasing, the first inequality (7a) follows. Inequality (7b) holds because x is decharged. The last inequality (7c) follows from (6). We conclude for resource a

$$\begin{aligned} c_a(a, x_{-j}) &\geq \xi_j(a, x_{-j}) \\ &\geq \xi_j(x) && (8) \\ &= \frac{c_{x_j}(x)}{\sum_{k\in S_{x_j}(x)} \min_{b\in M} c_b(b, x_{-k})} \cdot \min_{b\in M} c_b(b, x_{-j}) && (9) \\ &> 0 && (10) \\ &= c_a(z), && (11) \end{aligned}$$

where (8) holds because x is a pure Nash equilibrium and (9) stems from the definition of the protocol because there are no strong foreign players and hence $c_{x_j}(x) > 0$. Inequality (10) holds because of (7) and finally (11) holds because of (4). Hence, there must be a nonempty set of players $S_a(x) \backslash S_a(z)$. These players cannot be strong foreign players, thus $c_a(x) > 0$. With $c_a(z) = 0$ and $c_a(x) > 0$ we have $\ell_a(x) > \ell_a(z)$ for all resources $a \neq z_i$ with foreign players. For all resources a without foreign players we know $\ell_a(x) \geq \ell_a(z)$ and for resource x_i even $\ell_{x_i}(x) = \ell_{x_i}(z) + d_i$ because we are in Case (i). Because the total load is the same in x and z, we have for resource z_i

$$\ell_{z_i}(z_i, x_{-i}) = \ell_{z_i}(x) + d_i \leq \ell_{z_i}(z). \tag{12}$$

Consequently,

$$\begin{aligned} \xi_i(z) = c_{z_i}(z) &\geq c_{z_i}(z_i, x_{-i}) && (13) \\ &\geq \frac{c_{x_i}(x)}{\sum_{j\in S_{x_i}(x)} \min_{b\in M} c_b(b, x_{-j})} \cdot \min_{b\in M} c_b(b, x_{-i}) && (14) \\ &= \xi_i(x) = \xi_i(x_i, z_{-i}), && (15) \end{aligned}$$

where the first inequality (13) holds because of (12) and the second inequality (14) because x is decharged and $c_{z_i}(z_i, x_{-i}) \geq \min_{b\in M} c_b(b, x_{-i})$. Equality (15) holds by the definition of the strong x-enforcing protocol for Case (i) and the last equation holds because we assume Case (i). If $|S_{x_i}(x)| > 1$, then inequality (14) is strict, because x is strongly decharged (i.e., (2) holds), which implies $\xi_i(z) > \xi_i(x_i, z_{-i})$. This contradicts the fact that z is a pure Nash equilibrium. Thus, $S_{x_i}(x) = \{i\}$ and $c_{z_i}(z) = \xi_i(z) \leq \xi_i(x_i, z_{-i}) = c_{x_i}(x_i, z_{-i}) = c_{x_i}(x)$. Moreover, using $c_{x_i}(z) = 0$, because $\ell_{x_i}(z) = 0$, we obtain $c_{x_i}(z) + c_{z_i}(z) \leq c_{x_i}(x) + c_{z_i}(x)$ as desired.

Case (ii) is $S_{x_i}(x_i, z_{-i}) \subset S_{x_i}(x)$ and $i = \min S_{x_i}(x_i, z_{-i})$. Here, we obtain

$$c_{z_i}(z) = \xi_i(z) \le \xi_i(x_i, z_{-i}) = c_{x_i}(x_i, z_{-i}) \le c_{x_i}(x),$$

where the first inequality holds because z is a pure Nash equilibrium. The second inequality holds because Case (ii) implies $\ell_{x_i}(x_i, z_{-j}) \le \ell_{x_i}(x)$. We also have

$$c_{x_i}(z) = \sum_{j \in S_{x_i}(z)} \xi_j(z) \le \sum_{j \in S_{x_i}(z)} \xi_j(z_i, z_{-j}) = 0.$$

This inequality holds because z is Nash equilibrium and because in this case all players $j \in S_{x_i}(z)$ have a higher index $j > i$. Consequently, we have again $c_{x_i}(z) + c_{z_i}(z) \le c_{x_i}(x) + c_{z_i}(x)$.

Finally, we examine Case (iii) where the protocol returns $\xi_i(x_i, z_{-i}) = 0$ and thus for the pure Nash equilibrium z we have $c_{z_i}(z) = \xi_i(z) \le \xi_i(x_i, z_{-i}) = 0$. Again, $c_{x_i}(z) + c_{z_i}(z) \le c_{x_i}(x) + c_{z_i}(x)$. □

Remark 3.1. In contrast to the characterization of Theorem 3.1, the most expensive pure Nash equilibrium of a game is not necessarily strongly decharged. As an example consider the model with two players, $d_1 = d_2 = 1$ and two resources with identical cost functions satisfying $c_{a_1}(1) = c_{a_2}(1) = 1$ and $c_{a_1}(2) = c_{a_2}(2) = 2$. Here, all profiles have the same cost and $x = (a_1, a_1)$ is a pure Nash equilibrium under the x-enforcing protocol, but x is not strongly decharged.

3.2.1. An optimal protocol. Using the insights gained in the previous sections, we show that among all basic and separable protocols, the x-enforcing protocol gives rise to an optimal protocol simultaneously minimizing the price of anarchy and stability. Our main result involves the nth harmonic number $\mathscr{H}_n = \sum_{i=1}^{n} (1/i)$.

Theorem 3.3. *For singleton games,*

$$\min_{\Xi \in \mathscr{B}_n} \mathrm{PoS}(\Xi) = \min_{\Xi \in \mathscr{B}_n} \mathrm{PoA}(\Xi) = \min_{\Xi \in \mathscr{S}_n} \mathrm{PoS}(\Xi) = \min_{\Xi \in \mathscr{S}_n} \mathrm{PoA}(\Xi) = \mathscr{H}_n.$$

We prove the theorem by two subsequent lemmas. We first present an algorithm that returns for any singleton model a strongly decharged profile of cost at most $\mathscr{H}_n$ times the cost of an optimal profile. Hence, we can design a protocol that for any singleton model

(i) runs the algorithm to obtain a strongly decharged profile x, and
(ii) uses x as input for the x-enforcing protocol to obtain cost share functions.

This protocol is separable and has PoA as well as PoS of $\mathscr{H}_n$. In the second lemma, we show that the protocol is optimal by proving that $\mathscr{H}_n$ is a lower bound on the price of stability for every basic protocol.

Lemma 3.3. *Any singleton model (N, M, d, c) with an optimal profile y has a strongly decharged profile x with $C(x) \le \mathscr{H}_n \cdot C(y) = \sum_{k=1}^{n} (1/k) \cdot C(y)$.*

Proof. The desired profile x is found by Algorithm 1. The algorithm takes as input an optimal profile y. In each cycle k of the algorithm's main loop (lines 4–20), a player i^k on the most expensive charged resource a^k is selected (line 5) and moved to the cheapest available resource b^k (lines 16, 17). If possible, the algorithm selects a player who can be moved to a cost-free resource, this is called *zero move* (line 6). Otherwise, it selects a player that has been moved before in a last-in/first-out scheme that is maintained through the stacks Q_a. Such moves are called *shuffles* (line 9). If neither a zero move nor a shuffle is possible, the smallest player on the resource is selected, which is called *kickoff* (line 12). The algorithm terminates when no charged resources are left. We show in two claims that the algorithm indeed terminates and that returned profiles are as cheap as desired.

Claim. Algorithm 1 *terminates.* Observe that shuffles are only performed when zero moves are not possible. Hence, if in cycle k a shuffle is performed, the following inequalities hold:

$$\min_{b \in M} c_b(b, x^k_{-j}) > 0 \quad \text{for all } j \in S_{a^k}(x^k). \tag{16}$$

We now consider two cases. For $|S_{a^k}(x^k)| = 1$, we obtain

$$c_{a^k}(x^k) > \min_{b \in M} c_b(b, x^k_{-i^k}) = c_{b^k}(b^k, x^k_{-i^k}) = c_{b^k}(x^{k+1}),$$

where the inequality follows because a^k is charged in x^k and the equality follows because Algorithm 1 moves i^k to the cheapest available resource.

Algorithm 1 (Find strongly decharged profile x)
Input: Singleton resource selection model (N, M, d, c), profile y
Output: Strongly decharged profile x

$k \leftarrow 1$ {stepnumber}
$x^1 \leftarrow y$ {starts with optimal profile y}
$Q_a \leftarrow \varnothing$ for all $a \in M$ {stacks that return the last element entered}
while there are charged resources **do**
  $a^k \leftarrow \arg\max \{c_a(x^k)\colon a \in M$ is charged$\}$ {select the most expensive charged resource}
  if $\min\{c_b(b, x^k_{-i})\colon b \in M\} = 0$ for $i = \min S_{a^k}(x^k)$ **then**
    {player can move to cost-free resource, case called *zero move*}
    $i^k \leftarrow \min S_{a^k}(x^k)$ {select smallest player}
  else if $Q_{a^k} \neq \varnothing$ **then**
    {some player on a^k was moved before, case called *shuffle*}
    $i^k \leftarrow$ extract from Q_{a^k} {select last moved player}
  else
    {no foreign players on a^k, case called *kickoff*}
    $i^k \leftarrow \min S_{a^k}(x^k)$ {select smallest player}
  end if
  $b^k \leftarrow \arg\min \{c_b(b, x^k_{-i^k})\colon b \in M\}$ {select cheapest resource}
  $x^{k+1} \leftarrow (b^k, x^k_{-i^k})$ {move player}
  enter i^k to stack Q_{b^k}
  $k \leftarrow k + 1$ {iterate}
end while
return $x \leftarrow x^k$

If $|S_{a^k}(x^k)| > 1$, then we obtain

$$c_{a^k}(x^k) \geq \sum_{j \in S_{a_k}(x^k)} \min_{b \in M} c_b(b, x^k_{-j}) \tag{17}$$

$$> \min_{b \in M} c_b(b, x^k_{-i^k}) = c_{b^k}(b^k, x^k_{-i^k}) = c_{b^k}(x^{k+1}), \tag{18}$$

where (17) is valid because a^k is charged in x^k. The second inequality (18) holds because of (16) and the equalities follow as above. In both cases, a shuffle moves the player to a strictly cheaper resource. To see that the algorithm terminates, we will now follow some player i over the course of the algorithm. Each zero move and each shuffle take her to a strictly cheaper resource. If the player is moved in cycle k and is next moved by a shuffle in cycle l, the cost of her resource $x_i^{k+1} = x_i^l$ may increase in the meantime as other players arrive on that resource. The algorithm assures by its last-in/first-out mechanism that these other players have been moved again before the shuffle in cycle l and consequently the cost has decreased to the original level $c_{x_i^{k+1}}(x^{k+1}) \geq c_{x_i^l}(x^l)$. Because only resources with positive costs can be charged, this implies that after a zero move, the player will never again be considered for shuffles. Hence, a player can be moved by at most one kickoff, afterward a sequence of shuffles and thereafter only zero moves. The sequence of shuffles is finite because each shuffle takes the player to a strictly cheaper resource. Once the player has been moved by a zero move, further zero moves are only possible if in between some other player arrives on the player's resource via a kickoff or a shuffle, but again this is only finitely often possible. Altogether, each player can only be moved finitely often and thus the algorithm terminates after a finite number of cycles.

Claim. The final profile x has cost $C(x) \leq \mathscr{H}_n \cdot C(y)$. The concept of this final part of the proof is that in profile x the cost of every used resource is determined by the player who has last moved there or, if there are no such players, the home players. For this, some new notation is needed. Let p_i, $i \in N$, correspond to the position (by index) of player i on her optimal resource y_i, i.e., on any resource a we have $p_j = 1$ for player $j = \max S_a(y)$, $p_{j'} = 2$ for $j' = \max(S_a(y) \backslash \{j\})$, and so on. Consequently, when some player i performs her kickoff in cycle k, there are p_i players sharing her resource $a^k = y_i$ at that moment and she is the smallest of them. We obtain for resource b^k that she is moved to

$$c_{b^k}(x^{k+1}) = c_{b^k}(b^k, x^k_{-i}) = \min_{b \in M} c_b(b, x^k_{-i})$$

$$\leq \frac{1}{p_i} \cdot \sum_{j \in S_{a^k}(x^k)} \min_{b \in M} c_b(b, x^k_{-j}) \tag{19}$$

$$\leq \frac{1}{p_i} \cdot c_{a^k}(x^k) \tag{20}$$

$$\leq \frac{1}{p_i} \cdot c_{y_i}(y), \tag{21}$$

where the first inequality (19) is valid because i is the smallest of the p_i players on resource a^k in step k, the second inequality (20) holds because a^k is charged in x^k and the last inequality (21) holds because there are no foreign players on $a^k = y_i$ and hence $\ell_{a^k}(x^k) \leq \ell_{a^k}(y) = \ell_{y_i}(y)$.

Because shuffles and zero moves assign player i to cheaper resources, after her last move in cycle k', she is on resource $b^{k'}$ at cost

$$c_{b^{k'}}(x^{k'}) \leq \frac{1}{p_i} \cdot c_{y_i}(y).$$

Altogether, in the final profile x, the cost of a resource $a \in M$ to which players have been moved is determined by the last player who was moved there, which we denote by i_a. We thus obtain $c_a(x) \leq (1/p_{i_a}) \cdot c_{y_{i_a}}(y)$. For resources $a \in M$ that are used in x but where no player has been moved, the player $i_a := \max S_a(y)$ with $p_{i_a} = 1$ is still on resource a. In this case, the cost is bounded from above by $c_a(x) \leq c_a(y) = (1/p_{i_a}) \cdot c_{y_{i_a}}(y)$. Unused resources $a \in M$ have cost $c_a(x) = 0$. Altogether, we obtain $c_a(x) = (1/p_{i_a}) \cdot c_{y_{i_a}}(y)$ for all $a \in M$ with $\ell_a(x) > 0$, and $c_a(x) = 0$ for all $a \in M$ with $\ell_a(x) = 0$. This yields the desired bound for the cost of profile x, because now every used resource $a \in M$ has a unique player i_a that determines the resource's cost. We obtain

$$C(x) = \sum_{a \in M} c_a(x) \leq \sum_{\substack{a \in M \\ \ell_a(x) > 0}} \frac{1}{p_{i_a}} \cdot c_{y_{i_a}}(y) \leq \sum_{i \in N} \frac{1}{p_i} \cdot c_{y_i}(y) \leq \sum_{a \in M} \mathcal{H}_{p_{\max}} \cdot c_a(y)$$

$$= \mathcal{H}_{p_{\max}} \cdot C(y) \leq \mathcal{H}_n \cdot C(y), \quad \text{where } p_{\max} := \max\{|S_a(y)| : a \in M\}.$$

Observe that the bound for the price of anarchy obtained here can be much lower than $\mathcal{H}_n$ for singleton models that have optimal profiles, where the players are scattered over the resources and where therefore $p_{\max}$ is smaller than n. □

Remark 3.2. Although Lemma 3.3 shows that an optimal profile can be turned into a strongly decharged profile of cost at most $\mathcal{H}_n$ times the cost of an optimal profile, this holds true more generally: Algorithm 2 turns every profile into a strongly decharged profile with a cost increase of a factor at most $\mathcal{H}_n$.

Whereas the previous Lemma showed how to find a strongly decharged profile of at most $\mathcal{H}_n$ times the cost of an optimal profile, we now present an instance that has no weakly decharged profile cheaper than $\mathcal{H}_n$ times the cost of an optimal profile, thus giving a lower bound on the PoS.

Lemma 3.4. *The price of stability is at least $\mathcal{H}_n$ for games induced by basic protocols. This lower bound holds even for models with unit demands.*

Proof. Consider the singleton model (N, M, d, c) with n players that have unit demand $d_i = 1$ for all $i \in N$ and n resources with cost functions as in Table 2.

The only optimal profile is clearly $y = (a_1, \ldots, a_1)$ with $C(y) = 1 + \epsilon$. A profile z can only be a pure Nash equilibrium if it is weakly decharged (Lemma 3.2). We show that the cheapest weakly decharged profiles are those in which each resource is used by exactly one player, which all have the same cost as $x = (a_1, \ldots, a_n)$. It is easy to see that profile x is weakly decharged and with $C(x) = \sum_{i=1}^{n}(1/i) = \mathcal{H}_n$ this proves the lemma.

If in a profile z some resource other than a_1 is used by multiple players, then $C(z) \geq n$, thus such profiles are more expensive than x. If in profile z multiple players use resource a_1, say k players, then there are at least

Table 2. Cost functions for resources used in the proof of Lemma 3.4.

ℓ	$c_{a_1}(\ell)$	$c_{a_2}(\ell)$	...	$c_{a_i}(\ell)$	...	$c_{a_n}(\ell)$
0	0	0	...	0	...	0
1	$1+\epsilon$	$\frac{1}{2}$	...	$\frac{1}{i}$	...	$\frac{1}{n}$
>1	$1+\epsilon$	n	...	n	...	n

Note. For some small $\epsilon > 0$.

Table 3. Cost functions of resources used in the proof of Theorem 3.4.

ℓ	$c_{a_1}(\ell)$	$c_{a_2}(\ell)$	$c_{a_3}(\ell)$	$c_{a_4}(\ell)$	$c_{a_5}(\ell)$	$c_{a_6}(\ell)$	$c_{a_7}(\ell)$
0	0	0	0	0	0	0	0
2	0	0	0	0	0	k^2	0
3	0	0	0	k^3	0	k^4	0
4	1	k	1	$2k^3$	k^3	$2k^4$	k^4
5	2	k^3	k^5		$2k^3$		
6	k^3	k^5					
7	k^4						

$k-1$ unused resources and for the cheapest of these, say resource $\hat{a}$, we have $c_{\hat{a}}(1) \leq 1/k$. Thus, z is not weakly decharged as

$$c_{a_1}(z) = 1+\epsilon > 1 = \sum_{i\in S_{a_1}(z)} \frac{1}{k} \geq \sum_{i\in S_{a_1}(z)} c_{\hat{a}}(\hat{a}, z_{-i}) = \sum_{i\in S_{a_1}(z)} \min_{b\in M} c_b(b, z_{-i}).$$

Altogether, only such profiles in which all resources are used by exactly one player are cheap weakly decharged profiles and x is the cheapest pure Nash equilibrium. □

3.3. Uniform protocols. The basic and separable protocols that we introduced so far were always tailored to a desirable (decharged) profile. Because uniform protocols need to work independent of the set M, they cannot be based on specific profiles. We show in this section that uniformity leads to an unbounded price of anarchy. The question of $\min_{\xi\in\mathcal{U}_n} \mathrm{PoS}(\xi)$ remains open.

Theorem 3.4. *There is no uniform protocol for which the price of anarchy has an upper bound. This holds even for models with at most three players, three resources, and nondecreasing costs per unit.*

Proof. The essence of uniform protocols is that adding resources to or removing them from the model does not change the cost shares of players using a certain resource, as long as the player set and the weight vector remain the same. This motivates the definition of cost share functions $\hat{\xi}_i$ that return the cost share ξ_i of player i as a function of the resource a that she uses and the set of players $S \subseteq N$ sharing the resource:

$$\hat{\xi}_i(a, S) := \xi_i(x) \quad \forall\, a\in M,\;\; S\subseteq N,\;\; i\in S,\;\; x\in M^n\colon S_a(x) = S.$$

As in Definition 2.1, a profile x is a pure Nash equilibrium if none of the players can reduce their private cost by choosing a different resource. This can be expressed via cost share functions as follows. For all $i\in N$, $a\in M$ it holds that

$$\hat{\xi}_i(x_i, S_{x_i}(x)) \leq \hat{\xi}_i(a, S_a(x)\cup\{i\}). \tag{22}$$

For the proof of the theorem, we propose a number of singleton models and show that for any uniform protocol at least one of these models has a pure Nash equilibrium of more than k times the cost of an optimal profile for arbitrary $k\geq 2$. Throughout the entire proof, the player set will always be $N=\{1,2,3\}$ with weights $d=(4,3,2)$. The resources will be a subset of $M=\{a_1,\ldots,a_7\}$ with cost functions as outlined in Table 3.

First, consider the model with resources $M_1=\{a_1,a_2\}$ and their respective cost functions. The optimal profile $y_1=(a_2,a_1,a_1)$ has cost $C(y_1)=k+2$, while the profile $x_1=(a_1,a_2,a_2)$ has cost $C(x_1)=k^3+1$. Either x_1 is a pure Nash equilibrium and hence the protocol has a price of anarchy greater than k or one of the three players can reduce her cost share by choosing a different resource, which results by (22) in the following three cases:

(a) $\hat{\xi}_1(a_2,\{1,2,3\}) < \hat{\xi}_1(a_1,\{1\}) = 1$. In this case, consider $M_2=\{a_2,a_3\}$. The optimal profile $y_2 = (a_3,a_2,a_2)$ with $C(y_2)=k^3+1$ is not a pure Nash equilibrium and because of stability some other profile x with cost $C(x)\geq k^5$ has to be a pure Nash equilibrium.

(b) $\hat{\xi}_2(a_1,\{1,2\}) < \hat{\xi}_2(a_2,\{2,3\}) \leq k^3$. In this case, consider $M_3=\{a_1,a_2,a_4\}$. The optimal profile $y_3 = (a_1,a_2,a_4)$ has cost $C(y_3)=1$, and the profile $x_3=(a_2,a_1,a_4)$ has cost $C(x_3)=k$. Either x_3 is a pure Nash equilibrium or, again, one of the players can reduce her cost share by choosing a different resource, which leads to the following cases:

(b.1) $\hat{\xi}_1(a_1,\{1,2\}) < \hat{\xi}_1(a_2,\{1\}) = k$. This contradicts (b), that is
$\hat{\xi}_1(a_1,\{1,2\}) = c_1(d_1+d_2) - \hat{\xi}_2(a_1,\{1,2\}) > k^4 - k^3 > k$.

(b.2) $\hat{\xi}_1(a_4,\{1,3\}) < \hat{\xi}_1(a_1,\{1\}) = k$. In this case, consider $M_4=\{a_2,a_4,a_5\}$. The optimal profile $y_4 = (a_2,a_5,a_4)$ with $C(y_4)=k$ is not a pure Nash equilibrium and because of stability some other profile k with cost $C(x)\geq k^3$ has to be a pure Nash equilibrium.

(b.3) Players 2 and 3 cannot reduce their cost share as $\xi_2(x_3) = c_{a_1}(x_3) = 0$ and $\xi_3(x_3) = c_{a_4}(x_3) = 0$.

(c) $\hat{\xi}_3(a_1, \{1,3\}) < \hat{\xi}_3(a_2, \{2,3\}) \le k^3$. In this case, consider $M_5 = \{a_1, a_2, a_6\}$. The optimal profile $y_5 = (a_2, a_1, a_1)$ has cost $C(y_5) = k+1$, and the profile $x_5 = (a_1, a_2, a_6)$ has cost $C(x_5) = k^2 + 1$. Either x_5 is a pure Nash equilibrium or, again, one of the players can reduce her cost share by choosing a different resource.

(c.1) $\hat{\xi}_1(a_2, \{1,2\}) < \hat{\xi}_1(a_1, \{1\}) = 1$. In this case, consider again $M_3 = \{a_1, a_2, a_4\}$. The optimal profile $y_3 = (a_1, a_2, a_4)$ with $C(y_3) = 1$ is not a pure Nash equilibrium and because of stability some other profile x with cost $C(x) \ge k$ has to be a pure Nash equilibrium.

(c.2) $\hat{\xi}_1(a_6, \{1,3\}) < \hat{\xi}_1(a_1, \{1\}) = 1$. In this case, consider $M_6 = \{a_3, a_5, a_6\}$. The optimal profile $y_6 = (a_3, a_5, a_6)$ with $C(y_6) = k^2 + 1$ is not a pure Nash equilibrium and because of stability some other profile x with cost $C(x) \ge k^3$ has to be a pure Nash equilibrium.

(c.3) Player 2 cannot reduce her cost share because $\xi_2(x_5) = c_{a_2}(x_5) = 0$.

(c.4) $\hat{\xi}_3(a_1, \{1,3\}) < \hat{\xi}_3(a_6, \{3\}) = k^2$. In this case, consider $M_7 = \{a_1, a_6, a_7\}$. The optimal profile $y_7 = (a_1, a_7, a_6)$ with $C(y_7) = k^2 + 1$ is not a pure Nash equilibrium and because of stability some other profile x with cost $C(x) \ge k^3$ has to be a pure Nash equilibrium.

(c.5) $\hat{\xi}_3(a_2, \{2,3\}) < \hat{\xi}_3(a_6, \{3\}) = k^2$. This extends the original assumption from (c), which is $\hat{\xi}_3(a_1, \{1,3\}) < \hat{\xi}_3(a_2, \{2,3\}) < k^2$. Therefore this case implies (c.4).

Altogether, every uniform protocol allows in at least one of the analyzed cases a pure Nash equilibrium of at least k times the cost of an optimal profile for an arbitrarily large $k \ge 2$. Consequently, the price of anarchy is not bounded. □

4. Matroid games. We now turn to matroid games, where each player chooses multiple resources that form a basis of some matroid (a definition will be given below). The singleton games analyzed so far are a subclass of matroid games and consequently all lower bounds also hold in this new setting. Notably, there is no bound on the price of anarchy for matroid games induced by uniform protocols. We find a characterization of pure Nash equilibria in games induced by separable protocols and use it to show that the PoS for these games is exactly $\mathscr{H}_n$. For the PoA, however, we find a structural difference when going from singleton to matroids, showing that the PoA is exactly n. Moreover, our characterization of pure Nash equilibria induced by separable protocols cannot be carried over to basic protocols. Nevertheless, we use insights gained from the characterization to design an optimal basic protocol with PoS and PoA equal to $\mathscr{H}_n$. Before we present the details, we recall some basic definitions of matroids.

4.1. Matroids. A resource selection model (N, M, Σ, d, c) is called *matroid model* if there is a matroid $\mathscr{M} = (M, \mathscr{I})$ such that the players' strategy sets Σ_i equal the set of bases of $\mathscr{M}$ for every $i \in N$. Recall that a nonempty anti-chain[1] $\mathscr{B} \subseteq 2^M$ is the base set of a matroid $\mathscr{M} = (M, \mathscr{I})$ on resource (*ground*) set M if and only if the following *basis exchange property* is satisfied: whenever $X, Y \in \mathscr{B}$, and $x \in X \setminus Y$, then there exists some $y \in Y \setminus X$ such that $X \setminus \{x\} \cup \{y\} \in \mathscr{B}$. For a comprehensive overview on matroid theory, the reader is referred to Schrijver [54]. To shorten notation, we define $X + a := X \cup \{a\}$ and $X - a := X \setminus \{a\}$ for each set $X \subseteq M$ and each element $a \in M$. We denote by $\Sigma_i = \mathscr{B}(\mathscr{M})$ the collection of *bases* (i.e., the inclusion-wise maximal independent sets) of matroid $\mathscr{M}$. Every basis of a matroid has the same number of elements. The number of elements in a basis is called the rank of $\mathscr{M}$, or simply $\mathrm{rk}(\mathscr{M})$. Hence, in a given matroid model, the players' strategies consist of $\mathrm{rk}(\mathscr{M})$ resources. We will sometimes adopt the interpretation that each player has $\mathrm{rk}(\mathscr{M})$ jobs that she "schedules" on certain machines in M. These schedules must of course correspond to a basis of the matroid.

4.2. Characterization of Nash equilibria for separable protocols. Throughout this section, contrary to §3.1, the players are assumed to be ordered by nonincreasing weights: $d_1 \ge d_2 \ge \cdots \ge d_n$. We extend the concept of decharged profiles to prepare a characterization of pure Nash equilibria.

Definition 4.1 (Decharged Profile). In a matroid model (N, M, Σ, d, c), a profile $X \in \Sigma$ is called weakly decharged, if

$$c_a(X) \le \sum_{i \in S_a(X)} \min_{\substack{b \in M \\ X_i + b - a \in \Sigma_i}} c_b(X_i + b - a, X_{-i}) \quad \text{for all } a \in M.$$

The profile is called strongly decharged if additionally

$$c_a(X) < \sum_{i \in S_a(X)} \min_{\substack{b \in M \\ X_i + b - a \in \Sigma_i}} c_b(X_i + b - a, X_{-i}) \quad \text{for all } a \in M\colon c_a(X) > 0 \text{ and } |S_a(X)| > 1.$$

Otherwise the profile is called charged.

We further introduce the simple X-enforcing protocol.

Definition 4.2 (Simple X-Enforcing Protocol). The protocol takes as input a weakly decharged pro-

[1] $\mathscr{B} \subseteq 2^M$ is an *anti-chain* (with respect to $(2^M, \subseteq)$) if $B, B' \in \mathscr{B}$, $B \subseteq B'$ implies $B = B'$.

file X. We use X to define for any profile Z and resource a the set of *foreign players* $S_a^1(Z) := S_a(Z)\backslash S_a(X)$. Then, the simple X-enforcing protocol assigns for all $i \in N$, $a \in M$ and $Z \in \Sigma$ the cost share functions

$$\xi_{i,a}(Z) := \begin{cases} \dfrac{\min_{b\in M:\, X_i+b-a\in\Sigma_i} c_b(X_i+b-a, X_{-i})}{\sum_{j\in S_a(X)} \min_{b\in M:\, X_j+b-a\in\Sigma_j} c_b(X_j+b-a, X_{-j})} \cdot c_a(X), & \text{if } S_a(Z)=S_a(X) \text{ and } c_a(X)>0,\\ c_a(Z), & \text{if } S_a^1(Z)\neq\varnothing \text{ and } i=\min S_a^1(Z),\\ c_a(Z), & \text{if } S_a^1(Z)=\varnothing,\ S_a(Z)\subset S_a(X) \\ & \quad \text{and } i=\min S_a(Z),\\ 0, & \text{else.}\end{cases}$$

Note that this protocol is an extension of the x-enforcing protocol from §3, but without strong foreign players.

THEOREM 4.1. *Consider a matroid model (N, M, Σ, d, c). A profile $X \in \Sigma$ is weakly decharged if and only if it is a pure Nash equilibrium in the game induced by some separable protocol.*

PROOF. We start with the "if" direction. If X is a pure Nash equilibrium under some separable protocol Ξ that assigns cost share functions ξ, then for any $i \in N$, and any $a^* \in X_i$, $b \notin X_i$ such that $X_i + b - a^* \in \Sigma_i$, we have

$$\sum_{a\in X_i} \xi_{i,a}(X) = \xi_i(X) \le \xi_i(X_i+b-a^*, X_{-i}) = \sum_{a\in X_i+b-a^*} \xi_{i,a}(X_i+b-a^*, X_{-i}) \tag{23}$$

$$= \sum_{a\in X_i\backslash\{a^*\}} \xi_{i,a}(X) + \xi_{i,b}(X_i+b-a^*, X_{-i}), \tag{24}$$

where inequality (23) follows because X is a pure Nash equilibrium and (24) follows because the protocol is separable. Consequently,

$$\xi_{i,a}(X) \le \min_{\substack{b\in M\\ X_i+b-a\in\Sigma_i}} \xi_{i,b}(X_i+b-a, X_{-i}) \quad \text{for all } a \in M,\ \ i \in S_a(X). \tag{25}$$

With

$$c_a(X) = \sum_{i\in S_a(X)} \xi_{i,a}(X) \le \sum_{i\in S_a(X)} \min_{\substack{b\in M\\ X_i+b-a\in\Sigma_i}} \xi_{i,b}(X_i+b-a, X_{-i}) \tag{26}$$

$$\le \sum_{i\in S_a(X)} \min_{\substack{b\in M\\ X_i+b-a\in\Sigma_i}} c_b(X_i+b-a, X_{-i}) \quad \text{for all } a\in M \tag{27}$$

we conclude that the profile X is weakly decharged. Here, (26) follows from (25) and (27) because the protocol is budget balanced.

We now turn to the "only if" direction of the proof. For a weakly decharged profile X, it is clear that the simple X-enforcing protocol is budget balanced, we only show that X is a pure Nash equilibrium. For all resources $a \in M$ with $c_a(X) > 0$ we are in the first case of the definition of the protocol, thus, we obtain

$$\begin{aligned}\xi_{i,a}(X) &= \frac{\min_{b\in M:\, X_i+b-a\in\Sigma_i} c_b(X_i+b-a, X_{-i})}{\sum_{j\in S_a(X)} \min_{b\in M:\, X_j+b-a\in\Sigma_j} c_b(X_j+b-a, X_{-j})} \cdot c_a(X)\\ &\le \min_{b\in M:\, X_i+b-a\in\Sigma_i} c_b(X_i+b-a, X_{-i})\\ &\le \min_{b\in M\backslash\{a\}:\, X_i+b-a\in\Sigma_i} \xi_{i,b}(X_i+b-a, X_{-i}) \quad \text{for all } i\in S_a(X),\end{aligned} \tag{28}$$

where the first inequality holds because the profile x is weakly decharged. For all other resources $a \in M$, we have $\xi_{i,a}(X) = c_a(X) = 0$ for all $i \in S_a(X)$. Thus, so far we have shown that no player can improve by exchanging a single resource $a \in X_i$ with another resource $b \notin X_i$, that is, playing a strategy $X_i + b - a$. We call such exchanges of single resources $(1, 1)$-*exchanges*. Using the $(1, 1)$-exchange property of bases of matroids, we now show that X is a pure Nash equilibrium, i.e., that $\xi_i(X) \le \xi_i(Z_i, X_{-i})$ for all $Z_i \in \Sigma_i$. This approach has been used before by Ackermann et al. [2] in the context of classic weighted matroid congestion games.

To this end, fix a $Z_i \in \Sigma_i$ and denote by $G(X_i \triangle Z_i)$ the bipartite graph (V, E) with $V := (X_i\backslash Z_i) \cup (Z_i\backslash X_i)$ and $E := \{(a, b)\colon a \in X_i\backslash Z_i, b \in Z_i\backslash X_i, (X_i + b - a) \in \Sigma_i\}$.

PROPOSITION 4.1 (SCHRIJVER [54]). *There exists a perfect matching in the graph $G(X_i \triangle Z_i)$.*

Consider such a matching and observe that for any matching edge (a, b) with $a \in X_i$, $b \in Z_i$ we have

$$\xi_{i,b}(Z_i, X_{-i}) = c_b(Z_i, X_{-i}) = c_b(X_i+b-a, X_{-i}) = \xi_{i,b}(X_i+b-a, X_{-i}) \tag{29a}$$

$$\ge \xi_{i,a}(X), \tag{29b}$$

where the equations in (29a) are given by the definition of the protocol and inequality (29b) follows from (28).

Because this holds for any matching edge and the matching is perfect, we conclude

$$\begin{aligned}\xi_i(X) &= \sum_{a\in X_i}\xi_{i,a}(X) = \sum_{a\in X_i\cap Z_i}\xi_{i,a}(X) + \sum_{a\in X_i\setminus Z_i}\xi_{i,a}(X)\\ &\le \sum_{a\in X_i\cap Z_i}\xi_{i,a}(X) + \sum_{b\in Z_i\setminus X_i}\xi_{i,b}(Z_i, X_{-i}) \qquad (30)\\ &= \sum_{a\in X_i\cap Z_i}\xi_{i,a}(Z_i, X_{-i}) + \sum_{b\in Z_i\setminus X_i}\xi_{i,b}(Z_i, X_{-i}) = \xi_i(Z_i, X_{-i}). \qquad (31)\end{aligned}$$

Inequality (30) follows from (29) and (31) follows from the definition of the protocol. □

Whereas our characterization of pure Nash equilibria in singleton games works for both separable and basic protocols, in matroid games the characterization is limited to separable protocols. This is due to the nature of nonsingleton games: when the players' strategies consist of multiple resources, the question whether a certain profile can be a pure Nash equilibrium can generally not be answered by looking at the resources one by one. A special case is the class of matroid games induced by separable protocols: here, if a player can decrease his cost share by moving to a different strategy, he can also decrease his cost share by a $(1,1)$-exchange (because of matroid properties) and this exchange of one resource in his strategy will not affect his cost share on the other resources (because of the separability of the protocol). Consequently, only in this unique setting we can again characterize pure Nash equilibria by a "per resource" condition.

Remark 4.1. The above characterization allows the same extension as shown for the characterization of pure Nash equilibria in singleton games in Corollary 3.1. It holds on a per instance basis and is independent of the social welfare function. Hence, the simple X-enforcing protocol is universally PoS-optimal among all separable protocols for any social welfare function.

4.3. Bounds for the price of stability for separable protocols. Just like for singleton games, our characterization of pure Nash equilibria in matroid games induced by a separable protocols allows us to design an optimal protocol.

Algorithm 2 (Find strongly decharged profile X)
Input: Matroid resource selection model (N, M, Σ, d, c), profile Y
Output: Strongly decharged profile X

```
k ← 1   {stepnumber}
X^1 ← Y   {starts with optimal profile Y}
Q_a ← ∅ for all a ∈ M   {stacks that return the last element entered}
while there are charged resources do
    a^k ← arg max{c_a(X^k): a ∈ M is charged}   {select the most expensive charged resource}
    if c_b(X_i^k + b − a^k, X_{−i}^k) = 0 for some i ∈ S_{a^k}(X^k) and some b ∈ M: X_i^k + b − a^k ∈ Σ_i then
        {player can have (1,1)-exchange to cost-free resource, case called zero move}
        i^k ← i   {select player this player}
    else if Q_{a^k} ≠ ∅ then
        {some player on a^k was moved before, case called shuffle}
        i^k ← extract from Q_{a^k}   {select last moved player}
    else
        {no foreign players on a^k, case called kickoff}
        i^k ← arg min_{i∈S_{a^k}(X^k)} min_{b∈M, X_{i^k}^k + b − a^k ∈ Σ_{i^k}} c_b(X_{i^k}^k + b − a^k, X_{−i^k}^k)   {select player with cheapest (1,1)-exchange}
    end if
    b^k ← arg min_{b∈M, X_{i^k}^k + b − a^k ∈ Σ_{i^k}} c_b(X_{i^k}^k + b − a^k, X_{−i^k}^k)   {select cheapest (1,1)-exchange}
    X^{k+1} ← (X_{i^k}^k + b^k − a^k, X_{−i}^k)   {execute (1,1)-exchange}
    enter i^k to stack Q_{b^k}
    k ← k + 1   {iterate}
end while
return X ← X^k
```

THEOREM 4.2. *For matroid games,*

$$\min_{\Xi\in\mathcal{S}_n} \mathrm{PoS}(\Xi) = \mathcal{H}_n.$$

PROOF. The lower bound of $\mathcal{H}_n$ on the PoS for separable protocols established in Lemma 3.4 for singleton games also holds for matroid games. Hence, we only need to prove the upper bound. To do this, we introduce an adaption of Algorithm 1 to matroid games that returns for any matroid model a strongly decharged profile of at most $\mathcal{H}_n$ times the cost of an optimal profile. To prove that Algorithm 2 works as desired, we follow the same steps as in the proof of Algorithm 1. First, we show that the algorithm terminates, then we deal with the cost of profile X.

Claim. Algorithm 2 terminates. For the proof, we adhere to the interpretation of players scheduling $\mathrm{rk}(\mathcal{M})$ jobs on the resources. We fix a player i and follow one of his jobs over the course of the algorithm, showing that it can only be moved finitely often. There can be at most one kickoff involving this job, afterward the job is always in the stack of the resource it is scheduled on. Say the algorithm is in cycle k and a shuffle involving our job is performed. Then,

$$\min_{\substack{b\in M\\ X_j^k+b-a^k\in\Sigma_j}} c_b(X_j^k+b-a^k, X_{-j}^k) > 0 \quad \text{for all } j\in S_{a^k}(X^k), \tag{32}$$

because otherwise the algorithm would perform a zero move. We now consider two cases. For $|S_{a^k}(X^k)|=1$, we obtain

$$c_{a^k}(X^k) > \min_{\substack{b\in M\\ X_{i^k}^k+b-a^k\in\Sigma_{i^k}}} c_b(X_{i^k}^k+b-a^k, X_{-i^k}^k) \tag{33}$$

$$= c_{b^k}(X_{i^k}^k+b-a^k, X_{-i^k}^k) = c_{b^k}(X^{k+1}), \tag{34}$$

where (33) follows because a^k is charged in X^k. Equality (34) follows because Algorithm 2 moves i^k to the cheapest available resource.

If $|S_{a^k}(X^k)|>1$, then we obtain

$$c_{a^k}(X^k) \geq \sum_{j\in S_a(X^k)} \min_{\substack{b\in M\\ X_j^k+b-a^k\in\Sigma_j}} c_b(X_j^k+b-a^k, X_{-j}) \tag{35}$$

$$> \min_{\substack{b\in M\\ X_i^k+b-a^k\in\Sigma_i}} c_b(X_i^k+b-a^k, X_{-i}) \tag{36}$$

$$= c_{b^k}(X_i^k+b^k-a^k, X_{-i}) = c_{b^k}(X^{k+1}),$$

where (35) is valid because a^k is charged in X^k. The second inequality (36) holds because of (32) and the equalities follow as above. In both cases, a shuffle moves the player to a strictly cheaper resource and the last-in/first-out mechanism of the algorithm makes sure that in between two shuffles the cost of the resource does not increase more than it decreases. Hence, the number of shuffles is limited. Also, the number of zero moves is limited, as shown in the proof of Algorithm 1. Hence, each job is moved only finitely often and the algorithm terminates.

Claim. Profile X returned by Algorithm 2 has at most $\mathcal{H}_n$ times the cost of the input profile Y. Throughout the proof of this claim, we regard the set $Q := \{q_1, \ldots\}$ of all jobs instead of the players they belong to. We denote the player to which a job q belongs by $i(q)$ and the resource on which job q is scheduled in profile Y by $y(q)$. If k is the cycle of the algorithm in which q has its kickoff, we define $p(q) := |S_{y(q)}(X^k)|$ to be the number of players on $y(q)$ in X^k. Then,

$$c_{b^k}(X^{k+1}) = c_{b^k}(X_{i(q)}^k+b_k-y(q), X_{-i(q)}^k) = \min_{\substack{b\in M\\ X_{i(q)}^k+b-y(q)\in\Sigma_{i(q)}}} c_b(X_{i(q)}^k+b-y(q), X_{-i(q)}^k)$$

$$\leq \frac{1}{p(q)}\cdot \sum_{j\in S_{y(q)}(X^k)} \min_{\substack{b\in M\\ X_j^k+b-y(q)\in\Sigma_j}} c_b(X_j^k+b-y(q), X_{-j}^k) \tag{37}$$

$$\leq \frac{1}{p(q)}\cdot c_{y(q)}(X^k) \tag{38}$$

$$\leq \frac{1}{p(q)}\cdot c_{y(q)}(Y), \tag{39}$$

where the first inequality (37) is valid because $i(q)$ has the cheapest alternative among the $p(q)$ players on $y(q)$ in X^k, the second inequality (38) holds because $y(q)$ is charged in X^k and the last inequality (39) holds because there are no foreign players on $y(q)$ and hence $\ell_{y(q)}(X^k) \le \ell_{y(q)}(Y)$. Because shuffles take job q to strictly cheaper resources, we have for any cycle l in which q is moved, $c_{b^l}(X^{l+1}) \le (1/p(q)) \cdot c_{y(q)}(Y)$. Altogether, in the final profile X, the cost of a resource $b \in M$ to which jobs have been moved is determined by the last job that was moved there, which we denote by q_b. We have $c_b(X) \le (1/p_{q_b}) \cdot c_{a(q_b)}(Y)$. For resources $b \in M$ that are used in X but where no player has been moved, we have $\ell_b(X) \le \ell_b(Y)$ and hence $c_b(X) \le c_b(Y)$. On these resources, we pick an arbitrary job, call it q_b, and set $p_{q_b} = 1$ so that again $c_b(X) \le (1/p_{q_b}) \cdot c_{a(q_b)}(Y)$. We thus obtain

$$C(X) = \sum_{b \in M} c_b(X) \le \sum_{\substack{b \in M \\ S_b(X) \neq \varnothing}} \frac{1}{p_{q_b}} \cdot c_{a(q_b)}(Y) \le \sum_{\substack{q \in Q \\ p(q) \text{ defined}}} \frac{1}{p(q)} \cdot c_{y(q)}(Y) \le \sum_{a \in M} \mathscr{H}_{p_{\max}} \cdot c_a(Y)$$

$$= \mathscr{H}_{p_{\max}} \cdot C(Y) \le \mathscr{H}_n \cdot C(Y), \quad \text{where } p_{\max} := \max\{|S_a(y)|: a \in M\}. \quad \square$$

Remark 4.2. Again, as observed for singleton games in Remark 3.2, the algorithm turns every profile into a strongly decharged profile with a cost increase of a factor at most $\mathscr{H}_n$ and can, hence, take approximations of the optimal profile as input.

4.4. Bounds for the price of anarchy for separable protocols. A decharged profile X can be made a pure Nash equilibrium by the simple X-enforcing protocol, but we cannot enforce it as the game's most expensive pure Nash equilibrium. Instead, we find that across games induced by the X-enforcing protocol with a cheap decharged profile as input, the PoA is n.

Theorem 4.3. *For matroid games,*

$$\min_{\Xi \in \mathscr{S}_n} \text{PoA}(\Xi) = n.$$

We prove this result in two steps. In the upcoming lemma we give a matroid model that under any separable protocol has a pure Nash equilibrium of n times the cost of the optimal profile, resulting in our lower bound. Thereafter, we show that pure Nash equilibria in games induced by the simple X-enforcing protocol cost at most n times as much as the optimal profile, resulting in our upper bound.

Lemma 4.1. *For matroid games,*

$$\min_{\Xi \in \mathscr{S}_n} \text{PoA}(\Xi) \ge n.$$

This lower bound holds even for models with unit demands and uniform matroids.

Proof. Consider the matroid model (N, M, Σ, d, c) with n players that have unit demand $d_i = 1$ for all $i \in N$ and resources $M = \{b_0, b_1, \ldots, b_n\}$ with cost functions as in Table 4. All players $i \in N$ have identical strategy sets $\Sigma_i = \{Z_i \subset M: |Z_i| = n\}$. Note that the sets in Σ_i are the bases of the uniform matroid on M of rank n.

Consider the profile X with $X_i = \{b_1, \ldots, b_n\}$ for all players $i \in N$. We show for any separable protocol that assigns cost share functions $\xi_{i,a}$ that X is a pure Nash equilibrium. First, note that any strategy $Z_i \in \Sigma_i$, $Z_i \neq X_i$, can be written as $Z_i = X_i + b_0 - b_j$ for some $j \in N$. Now,

$$\xi_i(X) = \sum_{a \in X_i} \xi_{i,a}(X) \le \sum_{a \in X_i} \xi_{i,a}(X) - \xi_{i,b_j}(X) + \xi_{i,b_0}(X_i + b_0 - b_j, X_{-i}) \tag{40}$$

$$= \xi_i(X_i + b_0 - b_j, X_{-i}) = \xi_i(Z_i, X_{-i}), \tag{41}$$

where (40) holds because $\xi_{i,b_j}(X) \le c_{b_j}(X) = 1 = \xi_{i,b_0}(X_i + b_0 - b_j, X_{-i})$ because of budget balance of the protocol and (41) because $\xi_{i,a}(X) = \xi_{i,a}(X_i + b_0 - b_j, X_{-i})$ for all $a \in X_i$, $a \neq b_j$ because of separability of the protocol. Clearly, Y with $Y_i = M \setminus \{b_i\}$ is an optimal profile with $C(Y) = 1$. Hence, the fact that X with $C(X) = n$ is a pure Nash equilibrium under any separable protocol proves that the PoA is a least n. $\square$

Table 4. Cost functions for resources used in the proof of Lemma 4.1.

Load	$c_{b_0}(\ell)$	$c_{b_1}(\ell)$	...	$c_{b_n}(\ell)$
$\ell = 0$	0	0	...	0
$0 < \ell < n$	1	0	...	0
$\ell = n$	1	1	...	1

LEMMA 4.2. *For matroid games,*

$$\min_{\Xi\in\mathcal{S}_n} \mathrm{PoA}(\Xi) \le n.$$

PROOF. Let (N, M, Σ, d, c) be a matroid model with optimal profile Y. Let X be the strongly decharged profile returned by Algorithm 2 with intermediate profiles $X^1, \ldots$ and let $\xi_{i,a}$ for $i \in N$, $a \in M$ be the cost share functions assigned by the simple X-enforcing protocol. We use notation for players and jobs interchangeably, denoting the set of jobs on resource a by $q \in S_a(X)$. For each job q, define $y(q)$ and $p(q)$ as in the analysis of the algorithm (proof of Theorem 4.2) and denote additionally by $x(q)$ the resource $a \in M$ that job q is on in profile X and by X^q the algorithm's intermediate profile in which q is first on $x(q)$. From the analysis of the algorithm, we know

$$c_{x(q)}(X^q) \le \frac{1}{p(q)} c_{y(q)}(Y) \tag{42}$$

for all jobs q that were moved by the algorithm. For jobs q that were not moved by the algorithm, we set $X^q := Y$.

To prove the lemma, we show $C(Z) \le n \cdot C(Y)$ for any pure Nash equilibrium Z. To this end, we fix such a profile Z and link it to the profiles (X_i, Z_{-i}) via the Nash property,

$$C(Z) = \sum_{i\in N} \xi_i(Z) \le \sum_{i\in N} \xi_i(X_i, Z_{-i}) = \sum_{i\in N} \sum_{\substack{a\in X_i\\ S_a^1(Z)=\varnothing}} \xi_{i,a}(X_i, Z_{-i}) = \sum_{\substack{a\in M\\ S_a^1(Z)=\varnothing}} \sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i})$$

$$\le \sum_{\substack{a\in M\\ S_a^1(Z)=\varnothing}} \left(\sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + 2\cdot \sum_{q\in S_a(X)\setminus S_a(Y)} \frac{1}{p(q)} c_{y(q)}(Y) \right). \tag{43}$$

Proving (43) is a major challenge of this proof and beforehand we give a brief intuition for this inequality: for jobs that are moved by the algorithm we have an at most logarithmic cost increase going from profile Y to profile Z, represented by the second term, and for jobs not moved by the algorithm, the cost increase can even be linear as represented by the first term. In our worst-case example in Lemma 4.1, this linear cost increase dominates the logarithmic cost increase: no jobs are moved by the algorithm.

To prove (43), we partition the resources without foreign players into two sets:
- $M_1 := \{a \in M\colon S_a^1(Z) = \varnothing \text{ and } |S_a(X)\setminus S_a(Z)| \le 1\}$—resources, where at most one job is missing, and
- $M_2 := \{a \in M\colon S_a^1(Z) = \varnothing \text{ and } |S_a(X)\setminus S_a(Z)| > 1\}$—resources, where multiple jobs are missing.

For the resources in both sets, we find bounds corresponding to (43) in two separate claims. These claims combined prove the lemma.

CLAIM 4.3.1. *For $a \in M_1$,*

$$\sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i}) \le \sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + \sum_{q\in S_a(X)\setminus S_a(Y)} \frac{2}{p(q)} c_{y(q)}(Y).$$

PROOF. Recall that

$$c_a(X) \le \begin{cases} c_a(Y) & \text{for } a\in M \text{ with } S_a(X)\setminus S_a(Y) = \varnothing, & (44)\\[2mm] c_a(X^{q_a}) \le \dfrac{1}{p(q_a)} c_{y(q_a)}(Y) & \text{for } a \in M \text{ with } S_a(X)\setminus S_a(Y) \ne \varnothing, & (45)\end{cases}$$

where job q_a denotes the last job moved to a by the algorithm. The second inequality (45) follows from (42). To prove the claim, we have for $a \in M_1$,

$$\sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i}) = \mathbf{1}_{S_a(Z)\ne\varnothing} \sum_{i\in S_a(Z)} \xi_{i,a}(Z) + \mathbf{1}_{S_a(X)\setminus S_a(Z)=\{i^*\}} \xi_{i^*,a}(X_{i^*}, Z_{-i^*}) \tag{46}$$

$$\le \mathbf{1}_{S_a(Z)\ne\varnothing} c_a(X) + \mathbf{1}_{S_a(X)\setminus S_a(Z)=\{i^*\}} c_a(X) \tag{47}$$

$$\le \sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + 2\cdot \sum_{q\in S_a(X)\setminus S_a(Y)} \frac{1}{p(q)} c_{y(q)}(Y), \tag{48}$$

where $\mathbf{1}$ denotes the indicator function tied to the condition in subscript and, if applicable, i^* is the single player using resource a in X but not in Z. Equation (46) is due to the nature of M_1 and for (47) we use that there are no foreign players on machines $a \in M_1$. In inequality (48), we estimate using both (44) and (45). For the case where both indicator functions are true, we multiply the term from (45) by two. For the term from (44) this is not necessary because, if both indicator functions are true and $S_a(X)\backslash S_a(Y) = \varnothing$, then there are multiple jobs $q \in S_a(X) \cap S_a(Y)$ and, hence, $\sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) \geq 2 \cdot c_a(Y)$. □

Claim 4.3.2. *For $a \in M_2$,*

$$\sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i}) \leq \sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + \sum_{q\in S_a(X)\backslash S_a(Y)} \frac{1}{p(q)} c_{y(q)}(Y).$$

Proof. We denote the jobs on $a \in M_2$ in profile X by $q_1^a, \dots, q_{|S_a(X)|}^a$ such that they are indexed with nonincreasing weights $d_{q_1^a} \geq \cdots \geq d_{q^a_{|S_a(X)|}}$. Let $s(a) := \min\{i\colon q_i^a \in S_a(Z)\}$. Because the jobs are indexed in the same order as their players, the protocol assigns for $i \leq |S_a(X)|$

$$\xi_{q_i^a,a}(X_{q_i^a}, Z_{-q_i^a}) = \begin{cases} c_a(\ell_a(Z) + d_{q_i^a}) & \text{if } i < s(a), \\ c_a(\ell_a(Z)) & \text{if } i = s(a), \\ 0 & \text{if } i > s(a). \end{cases} \tag{49}$$

We now define an automorphism $\sigma_a\colon \{q_1^a, \dots, q_{|S_a(X)|}^a\} \to \{q_1^a, \dots, q_{|S_a(X)|}^a\}$ that maps the first $r(a) := |S_a(X)\backslash S_a(Y)|$ jobs (by index) to $S_a(X)\backslash S_a(Y)$, such that

- $\sigma_a(q_1^a)$ is the last job that was moved to a by the algorithm,
- $\sigma_a(q_2^a)$ is the second-last job that was moved to a by the algorithm,
- ... and
- $\sigma_a(q_{r(a)}^a)$ is the first job that was moved to a by the algorithm.

The remaining jobs are mapped arbitrarily to $S_a(X) \cap S_a(Y)$, keeping σ_a bijective. Then,

$$\ell_a(Z) \leq \ell_a(X) - \sum_{j=1}^{s(a)-1} d_{q_j^a} \tag{50a}$$

$$\leq \ell_a(X) - \sum_{j=1}^{s(a)-1} d_{\sigma_a(q_j^a)} \tag{50b}$$

$$\leq \ell_a(X^{\sigma_a(q_{s(a)}^a)}), \tag{50c}$$

where (50a) holds because $S_a(Z) \subset S_a(X)$ and $q_1^a, \dots, q_{s(a)-1}^a \notin S_a(Z)$ by definition of $s(a)$. Inequality (50b) holds because we indexed the jobs from big to small and hence the first $s(a)-1$ jobs are the "biggest" jobs on resource a. For (50c), if $s(a) \leq r(a)$, that is, if $\sigma_a(q_{s(a)}^a)$ was moved to resource a, then in profile $X^{\sigma_a(q_{s(a)}^a)}$ none of the jobs $\sigma_a(q_{s(a)-1}^a), \dots, \sigma_a(q_1^a)$ moved to a after job $\sigma_a(q_{s(a)}^a)$ are on resource a, and consequently (50c) follows. Otherwise, if $s(a) > r(a)$, that is, if $\sigma_a(q_{s(a)}^a) \in S_a(Y)$, then $Y = X^{\sigma_a(q_{s(a)}^a)}$ and in this profile none of the jobs $\sigma_a(q_{r(a)}^a), \dots, \sigma_a(q_1^a)$ that were moved to resource a are on resource a and hence (50c) follows. We find likewise for $i < s(a)$,

$$\ell_a(Z) + d_{q_i^a} \leq \ell_a(X) - \sum_{j=1}^{i-1} d_{q_j^a} \leq \ell_a(X) - \sum_{j=1}^{i-1} d_{\sigma_a(q_j^a)} \leq \ell_a(X^{\sigma_a(q_i^a)}), \tag{51}$$

where the above inequalities hold for similar reasons as (50). We complete the proof of Claim 4.3.2 by

$$\sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i}) = c_a(\ell_a(Z)) + \sum_{i=1}^{s(a)-1} c_a(\ell_a(Z) + d_{q_i^a}) \tag{52}$$

$$\leq \sum_{i=1}^{s(a)} c_a(\ell_a(X^{\sigma_a(q_i^a)})) \leq \sum_{q\in S_a(X)} c_a(X^{\sigma_a(q)}) \tag{53}$$

$$= \sum_{q\in S_a(X)} c_a(X^q) \tag{54}$$

$$\leq \sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + \sum_{q\in S_a(X)\backslash S_a(Y)} \frac{1}{p(q)} c_{y(q)}(Y), \tag{55}$$

where Equation (52) follows from (49) and inequality (53) follows from (50) and (51). Equation (54) holds because σ_a is an automorphism on $S_a(X)$ and finally inequality (55) follows from our definition of the intermediate profiles X^q and our results regarding these profiles as in (42). $\square$

We now continue the proof of Lemma 4.2 where we left off with (43) and conclude across both sets

$$
\begin{aligned}
C(Z) &\leq \sum_{\substack{a\in M\\ S_a^1(Z)=\varnothing}} \sum_{i\in S_a(X)} \xi_{i,a}(X_i, Z_{-i}) \\
&\leq \sum_{a\in M_1\cup M_2}\left(\sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + 2\cdot \sum_{q\in S_a(X)\setminus S_a(Y)} \frac{1}{p(q)} c_{y(q)}(Y)\right) && (56)\\
&\leq \sum_{a\in M}\left(\sum_{q\in S_a(X)\cap S_a(Y)} c_a(Y) + \sum_{q\in S_a(Y)\setminus S_a(X)} \frac{2}{p(q)} c_a(Y)\right) && (57)\\
&\leq \sum_{a\in M}\left(|S_a(Y)\cap S_a(X)|\cdot c_a(Y) + \sum_{p=|S_a(Y)\cap S_a(X)|+1}^{|S_a(Y)|} \frac{2}{p}\cdot c_a(Y)\right) && (58)\\
&\leq \sum_{a\in M} |S_a(Y)|\cdot c_a(Y) \leq n\cdot C(Y).
\end{aligned}
$$

Here, (56) follows from Claims 4.3.1 and 4.3.2. In (57), we change the order of summation: instead of summing up the $q\in S_a(X)\setminus S_a(Y)$ that were moved *from* other resources by the algorithm, we sum up the $q\in S_a(Y)\setminus S_a(X)$ that were moved *to* other resources by the algorithm. For (58), recall how we introduced $p(q)$: the first job that is moved away has $p(q)=|S_a(Y)|$, the next has $p(q)=|S_a(Y)-1|$ until the last job that is moved away has $p(q)=|S_a(Y)\cap S_a(X)|+1$. $\square$

4.5. An optimal basic protocol. For matroid games induced by basic protocols, we do not give a complete characterization of pure Nash equilibria. Instead, we find that the "decharged" condition is sufficient to make a profile a pure Nash equilibrium and even the most expensive pure Nash equilibrium of the game. Using this structural result, we give an optimal basic protocol that simultaneously minimizes PoS and PoA.

Lemma 4.3. *Consider a matroid model (N, M, Σ, d, c) and a profile $X\in\Sigma$. If X is strongly decharged, there is a basic protocol that assigns cost share functions such that X is the most expensive pure Nash equilibrium of the induced game.*

For the proof of this lemma, we introduce the anarchy eliminating protocol that just like the simple X-enforcing protocol relies on a decharged profile X.

Definition 4.3 (Anarchy Eliminating Protocol). The *anarchy eliminating protocol* takes as input a strongly decharged profile X. For any profile Z denote a global set of foreign players by $S^1(Z):=\{i\in N\colon Z_i\neq X_i\}$. Then, the anarchy eliminating protocol assigns for all $i\in N$ and $a\in M$ the cost share functions

$$
\xi_{i,a}(Z) := \begin{cases}
\dfrac{\min_{b\in M:\, X_i+b-a\in\Sigma_i} c_b(X_i+b-a, X_{-i})}{\sum_{j\in S_a(X)} \min_{b\in M:\, X_j+b-a\in\Sigma_i} c_b(X_j+b-a, X_{-j})}\cdot c_a(X), & \text{if } Z=X,\\
c_a(Z), & \text{if } Z\neq X,\ S^1(Z)\cap S_a(Z)\neq\varnothing \text{ and } i=\min S^1(Z)\cap S_a(Z),\\
c_a(Z), & \text{if } Z\neq X,\ S^1(Z)\cap S_a(Z)=\varnothing \text{ and } i=\min S_a(Z).\\
0, & \text{else.}
\end{cases}
$$

Proof of Lemma 4.3. We first show that the profile X is a pure Nash equilibrium under the anarchy eliminating protocol. To this end, let ξ^* be the cost share functions assigned by the simple X-enforcing protocol. Then, for any $i\in N$ and any $Z_i\in\Sigma_i$,

$$
\begin{aligned}
\xi_i(X) &= \sum_{a\in X_i} \xi_{i,a}(X) = \sum_{a\in X_i} \xi^*_{i,a}(X) \leq \sum_{a\in Z_i} \xi^*_{i,a}(Z_i, X_{-i}) && (59)\\
&\leq \sum_{a\in Z_i} c_a(Z_i, X_{-i}) = \sum_{a\in Z_i} \xi_{i,a}(Z_i, X_{-i}) = \xi_i(Z_i, X_{-i}), && (60)
\end{aligned}
$$

where (59) holds because X is a pure Nash equilibrium under the simple X-enforcing protocol and (60) holds because the simple X-enforcing protocol is budget balanced.

To complete the proof, we show that X is the most expensive pure Nash equilibrium, that is, $C(Z) \le C(X)$ for any pure Nash equilibrium Z. We denote the foreign player with the smallest index by $i^* := \min S^1(Z)$. Then,

$$\xi_j(Z) \le \xi_j(Z_{i^*}, Z_{-j}) = 0 \quad \text{for all } j \in S^1(Z), \quad j \ne i^*, \tag{61}$$

because Z is a pure Nash equilibrium. To estimate $\xi_{i^*}(Z)$, we need to examine three cases:

(i) $S^1(Z) = \{i^*\}$ and $S_a(X) = \{i^*\}$ for all $a \in X_{i^*}$,

(ii) $S^1(Z) = \{i^*\}$ and $|S_a(X)| > 1$ for some $a \in X_{i^*}$,

(iii) $|S^1(Z)| > 1$.

In Case (i), we have $c_a(Z) = c_a(X)$ for all $a \in M \setminus (Z_{i^*} \cup X_{i^*})$ and $S_a(Z) = \varnothing$ for all $a \in X_{i^*} \setminus Z_{i^*}$ and hence with

$$\begin{aligned} C(Z) &= \sum_{a \in M \setminus (Z_{i^*} \cup X_{i^*})} c_a(Z) + \sum_{a \in X_{i^*} \setminus Z_{i^*}} c_a(Z) + \sum_{a \in Z_{i^*}} c_a(Z) \\ &= \sum_{a \in M \setminus (Z_{i^*} \cup X_{i^*})} c_a(Z) + \sum_{a \in Z_{i^*}} \xi_{i^*,a}(Z) && (62) \\ &\le \sum_{a \in M \setminus (Z_{i^*} \cup X_{i^*})} c_a(X) + \sum_{a \in X_{i^*}} \xi_{i^*,a}(X_{i^*}, Z_{-i^*}) && (63) \\ &= \sum_{a \in M \setminus (Z_{i^*} \cup X_{i^*})} c_a(X) + \sum_{a \in X_{i^*}} c_a(X) = C(X) && (64) \end{aligned}$$

we are done. Equation (62) holds because $c_a(Z) = 0$ for all $a \in X_{i^*} \setminus Z_{i^*}$, inequality (63) holds because Z is a pure Nash equilibrium and (64) because $S_a(X) = \{i^*\}$ for all $a \in X_{i^*}$.

In Case (ii), because X is strongly decharged, we have $\xi_{i^*}(X) < \xi_{i^*}(Z_i, X_{-i^*}) = \xi_{i^*}(Z)$. This contradicts that Z is a pure Nash equilibrium.

In Case (iii), for resources that are used by players from $S^1(Z)$, we have

$$\begin{aligned} \sum_{\substack{a \in M \\ S_a(Z) \cap S^1(Z) \ne \varnothing}} c_a(Z) &= \sum_{i \in S^1(Z)} \xi_i(Z) && (65a) \\ &= \xi_{i^*}(Z) && (65b) \\ &\le \xi_{i^*}(X_{i^*}, Z_{-i^*}) && (65c) \\ &= \sum_{\substack{a \in X_{i^*} \\ i^* = \min S_a(X) \\ S_a(Z) \cap S^1(X_{i^*}, Z_{-i^*}) = \varnothing}} c_a(X_{i^*}, Z_{-i^*}) && (65d) \\ &\le \sum_{\substack{a \in X_{i^*} \\ i^* = \min S_a(X) \\ S_a(Z) \cap S^1(X_{i^*}, Z_{-i^*}) = \varnothing}} c_a(X), && (65e) \end{aligned}$$

where (65a) is given by the protocol, (65b) follows from (61), and (65c) holds because Z is a pure Nash equilibrium. Because there are multiple foreign players (Case (iii)), only the second case from the definition of the protocol applies for i^* and (65d) follows. Finally, (65e) holds because for resources $a \in X_{i^*}$ without foreign players, i.e., $S_a(Z) \cap S^1(X_{i^*}, Z_{-i^*}) = \varnothing$, we have $\ell_a(X_{i^*}, Z_{-i^*}) \le \ell_a(X)$.

We conclude,

$$\begin{aligned} C(Z) &= \sum_{\substack{a \in M \\ S_a(Z) \cap S^1(Z) = \varnothing \\ c_a(Z) > 0}} c_a(Z) + \sum_{\substack{a \in M \\ S_a(Z) \cap S^1(Z) \ne \varnothing}} c_a(Z) \\ &\le \sum_{\substack{a \in M \\ S_a(Z) \cap S^1(Z) = \varnothing \\ c_a(Z) > 0}} c_a(X) + \sum_{\substack{a \in X_{i^*} \\ i^* = \min S_a(X) \\ S_a(Z) \cap S^1(X_{i^*}, Z_{-i^*}) = \varnothing}} c_a(X) && (66) \\ &\le \sum_{a \in M} c_a(X) = C(X), && (67) \end{aligned}$$

where the first part of (66) holds because for resources $a \in M$ with $S_a(Z) \cap S^1(Z) = \varnothing$ we have $\ell_a(Z) \leq \ell_a(X)$ and the second part follows from (65). The two summands in (66) are disjoint, as for all resources $a \in X_{i^*}$ with $i^* = \min S_a(X)$ and $i^* \notin S_a(Z)$ we have $c_a(Z) = 0$, because $\xi_j(Z) \leq \xi_j(Z_{i^*}, Z_{-j}) = 0$ for all $j \in S_a(Z)$. Hence, (67) follows. □

Theorem 4.4. *For matroid games,*

$$\min_{\Xi \in \mathcal{B}_n} \mathrm{PoA}(\Xi) = \min_{\Xi \in \mathcal{B}_n} \mathrm{PoS}(\Xi) = \mathcal{H}_n.$$

Proof. When the anarchy eliminating protocol uses the strongly decharged profile generated by Algorithm 2 as input, it has PoS and PoA of $\mathcal{H}_n$. This matches the lower bound of $\mathcal{H}_n$ from Lemma 3.4.

5. Summary and discussion.

In this article we considered the design of cost sharing protocols for resource selection games with singleton and matroid structure. For the design goal of minimizing the resulting price of anarchy and price of stability, we obtained tight results for uniform, separable, and basic protocols.

5.1. Computational complexity.

All protocols introduced in this paper rely on optimal profiles and (strongly) decharged profiles. Hence, the complexity of computing such profiles is an essential question. For unweighted scheduling models, i.e., singleton models with $d_1 = \cdots = d_n$, our protocols are polynomial time computable. Chakrabarty et al. [14] have shown that one can compute an optimal profile in polynomial time by dynamic programming. Moreover, given an optimal profile, we can compute a (strongly) decharged profile in polynomial time using Algorithm 2. To estimate the runtime of Algorithm 2, note that for unweighted models the cost of a resource a depends only on the number of players using it. Using an idea of Ackermann et al. [1], we can write a list of all possible values $(c_a(\ell_a), \ell_a)$ that has length $m \cdot n$ and sort it lexicographically in nonincreasing order. We set markers to the list entries that correspond to the optimal profile and update the markers as the algorithm modifies the profile. In each cycle of Algorithm 2, a player is moved from the most expensive resource to a less expensive resource and, consequently, the position of the top marker is on a lexicographically lower list position. Considering the length of the list, there can be at most $m \cdot n$ cycles of the algorithm. It is easy to show that the worst-case runtime of the algorithm remains unchanged even for matroid models. We do not know, however, if one can compute the optimal profile in matroids in polynomial time. So far, tractability of computing an optimal profile is only known for matroid models with nondecreasing or nonincreasing costs per unit (cf. Ackermann et al. [1]).

For models with weighted demands the situation changes dramatically. Here, computing an optimal profile is NP-complete and not even approximable by any constant factor even for singleton models (see Hochbaum and Shmoys [34]). Even if an optimal profile is given, the runtime of Algorithm 2 appears to be superpolynomial. This raises two important questions: can we compute decharged profiles that satisfy our cost bounds and can we even directly compute the optimal decharged profile of a model in polynomial time?

5.2. Symmetry.

Although our results hold for general nondecreasing cost functions, we restricted the players' strategy sets to bases of a common matroid. Designing optimal protocols for more general strategy spaces while allowing for weighted players and nonmonotonic cost per units remains an important and practically relevant problem. We can give, however, a first insight showing that even basic cost sharing protocols perform significantly weaker when going from symmetric games to asymmetric games even in singleton games with unit demands.

Remark 5.1. In asymmetric singleton games, each player i has an individual strategy set $X_i \subseteq M$, opposed to $X_i = M$ for all $i \in N$ in symmetric games. The optimal bound on the PoA (Theorem 3.3) of $\mathcal{H}_n$ increases to at least n for asymmetric games. Consider the model (N, M, d, c) with n players, resources $M = \{a_0, a_1, \ldots, a_n\}$, unit weights $d_i = 1$, cost functions $c_a(\ell) = 1$ for $\ell \geq 1$, $a \in M$, and strategy sets $X_i = \{a_0, a_i\}$. The optimal profile is $y = (a_0, \ldots, a_0)$ with $C(y) = 1$, and $x = (a_1, \ldots, a_n)$ with $C(x) = n$ is a pure Nash equilibrium under every basic protocol.

5.3. Protocol design space.

Different applications allow for different means of cost sharing and, consequently, there is no generic choice of a suitable protocol design space. For this paper, we followed and extended the lead from Chen et al. [18] but other axioms defining the protocol space seem reasonable. For instance requiring the protocols to guarantee convergence of best-response dynamics or to ensure stability against deviations of coalitions of players seem equally natural.

Acknowledgments. The results of §3 appeared as an extended abstract in the proceedings of the 12th ACM Conference on Electronic Commerce (von Falkenhausen and Harks [56]). The authors thank the reviewers for helpful comments improving the presentation of the paper. The first author was supported by the Deutsche Forschungsgemeinschaft within the research training group 'Methods for Discrete Structures' (GRK 1408).

References

[1] Ackermann H, Röglin H, Vöcking B (2008) On the impact of combinatorial structure on congestion games. *J. ACM* 55(6):25:1–25:22.
[2] Ackermann H, Röglin H, Vöcking B (2009) Pure Nash equilibria in player-specific and weighted congestion games. *Theoret. Comput. Sci.* 410(17):1552–1563.
[3] Aland S, Dumrauf D, Gairing M, Monien B, Schoppmann F (2011) Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.* 40(5):1211–1233.
[4] Albers S (2009) On the value of coordination in network design. *SIAM J. Comput.* 38(6):2273–2302.
[5] Andrews M, Antonakopoulos S, Zhang L (2010) Minimum-cost network design with (dis)economies of scale. *Proc. 51st Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society, Washington, DC), 585–592.
[6] Anshelevich E, Dasgupta A, Kleinberg J, Tardos É, Wexler T, Roughgarden T (2008) The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38(4):1602–1623.
[7] Antonakopoulos S, Chekuri C, Shepherd FB, Zhang L (2011) Buy-at-bulk network design with protection. *Math. Oper. Res.* 36(1): 71–87.
[8] Archer A, Feigenbaum J, Krishnamurthy A, Sami R, Shenker S (2004) Approximation and collusion in multicast cost sharing. *Games Econom. Behav.* 47(1):36–71.
[9] Awerbuch B, Azar Y, Epstein A (2005) The price of routing unsplittable flow. Gabow HN, Fagin R, eds. *Proc. 37th Annual ACM Sympos. Theory Comput.* (ACM Press, New York), 57–66.
[10] Bhawalkar K, Gairing M, Roughgarden T (2010) Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. de Berg M, Meyer U, eds. *Proc. 18th Annual Euro. Sympos. on Algorithms, Part II*, Vol. 6347 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 17–28.
[11] Bilò V, Fanelli A, Flammini M, Melideo G, Moscardelli L (2010) Designing fast converging cost sharing methods for multicast transmissions. *Theory Comput. Systems* 47(2):507–530.
[12] Bogomolnaia A, Holzman R, Moulin H (2010) Sharing the cost of a capacity network. *Math. Oper. Res.* 35(1):173–192.
[13] Caragiannis I (2009) Efficient coordination mechanisms for unrelated machine scheduling. Mathieu C, ed. *Proc. 20th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 815–824.
[14] Chakrabarty D, Mehta A, Nagarajan V (2005) Fairness and optimality in congestion games. Riedl J, Kearns MJ, Reiter MK, eds. *Proc. 6th ACM Conf. Electronic Commerce* (ACM Press, New York), 52–57.
[15] Chekuri C, Chuzhoy J, Lewin-Eytan L, Naor J, Orda A (2006) Non-cooperative multicast and facility location games. Feigenbaum J, Chuang JC-I, Pennock DM, eds. *Proc. 7th ACM Conf. Electronic Commerce* (ACM Press, New York), 72–81.
[16] Chen H-L, Roughgarden T (2009) Network design with weighted players. *Theory Comput. Systems* 45(2):302–324.
[17] Chen Y-J, Zhang J (2012) Design of price mechanisms for network resource allocation via price of anarchy. *Math. Program.* 131(1–2):333–364.
[18] Chen H-L, Roughgarden T, Valiant G (2010) Designing network protocols for good equilibria. *SIAM J. Comput.* 39(5):1799–1832.
[19] Christodoulou G, Koutsoupias E (2005) On the price of anarchy and stability of correlated equilibria of linear congestion games. Brodal G, Leonardi S, eds. *Proc. 13th Annual European Sympos. Algorithms*, Number 3669 in Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 59–70.
[20] Christodoulou G, Koutsoupias E, Nanavati A (2004) Coordination mechanisms. Díaz J, Karhumäki J, Lepistö A, Sannella D, eds. *Proc. 31st Internat. Colloquium on Automata, Languages and Programming*, Volume 3142 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 45–56.
[21] Cole R, Correa JR, Gkatzelis V, Mirrokni VS, Olver N (2011) Inner product spaces for minsum coordination mechanisms. Fortnow L, Vadhan SP, eds. *Proc. 43rd Annual ACM Sympos. Theory Comput.* (ACM Press, New York), 539–548.
[22] Correa J, Schulz A, Stier-Moses N (2004) Selfish routing in capacitated networks. *Math. Oper. Res.* 29(4):961–976.
[23] Epstein A, Feldman M, Mansour Y (2009) Strong equilibrium in cost sharing connection games. *Games Econom. Behav.* 67(1):51–68.
[24] Even-Dar E, Kesselman A, Mansour Y (2007) Convergence time to Nash equilibrium in load balancing. *ACM Trans. Algorithms* 3(3):1–21.
[25] Fotakis D, Kontogiannis S, Spirakis P (2005) Selfish unsplittable flows. *Theoret. Comput. Sci.* (Springer, Berlin, Heidelberg), 348(2–3):226–239.
[26] Fotakis D, Kontogiannis S, Spirakis P (2006) Atomic congestion games among coalitions. Bugliesi M, Preneel B, Sassone V, Wegener I, eds. *Proc. 33rd Internat. Colloquium on Automata, Languages and Programming*, Vol. 4052 of Lecture Notes in Computer Science, 572–583.
[27] Fotakis D, Kontogiannis S, Koutsoupias E, Mavronicolas M, Spirakis P (2002) The structure and complexity of Nash equilibria for a selfish routing game. Widmayer P, Triguero F, Morales R, Hennessy M, Eidenbenz S, Conejo R, eds. *Proc. 29th Internat. Colloquium on Automata, Languages and Programming*, Vol. 2380 of Lecture Notes in Computer Science, 123–134.
[28] Gairing M, Monien B, Tiemann K (2006) Routing (un-) splittable flow in games with player-specific linear latency functions. Bugliesi M, Preneel B, Sassone V, Wegener I, eds. *Proc. 33rd Internat. Colloquium on Automata, Languages and Programming*, Vol. 4051 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 501–512.
[29] Goemans M, Mirrokni V, Vetta A (2005) Sink equilibria and convergence. *Proc. 46th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society, Washington, DC), 142–154.
[30] Harks T (2011) Stackelberg strategies and collusion in network games with splittable flow. *Theory Comput. Systems* 48(4):781–802.
[31] Harks T, Klimm M (2012) On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.* 37(3):419–436.

[32] Harks T, Klimm M, Möhring R (2009) Strong Nash equilibria in games with the lexicographical improvement property. Leonardi S, ed. *Proc. 5th Internat. Workshop on Internet and Network Econom.* Vol. 5929 of Lecture Notes in Computer Science, 463–470.
[33] Harks T, Miller K (2011) The worst-case efficiency of cost sharing methods in resource allocation games. *Oper. Res.* 59(6):1491–1503.
[34] Hochbaum DS, Shmoys DB (1987) Using dual approximation algorithms for scheduling problems theoretical and practical results. *J. ACM* 34(1):144–162.
[35] Hoefer M (2011) Competitive cost sharing with economies of scale. *Algorithmica* 60(4):743–765.
[36] Ieong S, McGrew R, Nudelman E, Shoham Y, Sun Q (2005) Fast and compact: A simple class of congestion games. Veloso MM, Kambhampati S, eds. *Proc. 20th National Conf. Artificial Intelligence and the 17th Innovative Appl. Artificial Intelligence Conf.* (AAAI Press/MIT Press, Cambridge, MA), 489–494.
[37] Immorlica N, Li L, Mirrokni V, Schulz A (2009) Coordination mechanisms for selfish scheduling. *Theoret. Comput. Sci.* 410(17):1589–1598.
[38] Johari R, Tsitsiklis J (2006) A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE J. Selected Area Comm.* 24(5):992–999.
[39] Johari R, Tsitsiklis JN (2009) Efficiency of scalar-parameterized mechanisms. *Oper. Res.* 57(4):823–839.
[40] Kollias K, Roughgarden T (2011) Restoring pure equilibria to weighted congestion games. Aceto L, Henzinger M, Sgall J, eds. *Proc. 38th Internat. Colloquium on Automata, Languages and Programming*, Number 6756 in Lecture Notes in Computer Science, 539–551.
[41] Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. Meinel C, Tison S, eds. *Proc. 16th Internat. Sympos. Theoret. Aspects of Comput. Sci.* Number 1563 in Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 404–413.
[42] Milchtaich I (1996) Congestion games with player-specific payoff functions. *Games Econom. Behav.* 13(1):111–124.
[43] Milchtaich I (2006) The equilibrium existence problem in finite network congestion games. Mavronicolas M, Kontogiannis S, eds. *Proc. 2nd Internat. Workshop on Internet and Network Econom.*, Vol. 4286 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 87–98.
[44] Moulin H (2002) Axiomatic cost and surplus sharing. Arrow KJ, Sen AK, Suzumura K, eds. *Handbook of Social Choice and Welfare*, 1st ed., Vol. 1, Chap. 06. (Elsevierl, North Holland), 289–357.
[45] Moulin H (2008) The price of anarchy of serial, average and incremental cost sharing. *Econom. Theory* 36(3):379–405.
[46] Moulin H (2010) An efficient and almost budget balanced cost sharing method. *Games Econom. Behav.* 70(1):107–131.
[47] Moulin H, Shenker S (2001) Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Econom. Theory* 18(3):511–533.
[48] Osborne M, Rubinstein A (1994) *A Course in Game Theory* (MIT Press, Cambridge, MA).
[49] Panagopoulou P, Spirakis P (2006) Algorithms for pure Nash equilibria in weighted congestion games. *ACM J. Exp. Algorithmics* 11(2.7):1–19.
[50] Rosenthal R (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2(1):65–67.
[51] Roughgarden T (2009) Intrinsic robustness of the price of anarchy. *Proc. 41st Annual ACM Sympos. Theory Comput.* (ACM Press, New York), 513–522.
[52] Roughgarden T, Tardos É (2002) How bad is selfish routing? *J. ACM* 49(2):236–259.
[53] Rozenfeld O, Tennenholtz M (2006) Strong and correlated strong equilibria in monotone congestion games. Mavronicolas M, Kontogiannis S, eds. *Proc. 2nd Internat. Workshop on Internet and Network Econom.*, Vol. 4286 of Lecture Notes in Computer Science (Springer, Berlin, Heidelberg), 74–86.
[54] Schrijver A (2003) *Combinatorial Optimization: Polyhedra and Efficiency* (Springer, Berlin).
[55] Srikant R (2003) *The Mathematics of Internet Congestion Control* (Birkhäuser, Boston).
[56] von Falkenhausen P, Harks T (2011) Optimal cost sharing protocols for scheduling games. Shoham Y, Chen Y, Roughgarden T, eds. *Proc. 12th ACM Conf. Electronic Commerce* (ACM Press, New York), 285–294.
[57] Yao F, Demers A, Shenker S (1995) A scheduling model for reduced CPU energy. *Proc. 35th Annual IEEE Sympos. Foundations Comput. Sci.* (IEEE Computer Society, Washington, DC), 374–382.