

Tractable mechanisms for computing near-optimal utility functions

Rahul Chandan
UC Santa Barbara
Santa Barbara, California
rchandan@ucsb.edu

Dario Paccagnan
Imperial College, London
London, United Kingdom
d.paccagnan@imperial.ac.uk

Jason R. Marden
UC Santa Barbara
Santa Barbara, California
jrmarden@ece.ucsb.edu

ABSTRACT

Large scale multiagent systems must rely on distributed decision making, as centralized coordination is either impractical or impossible. Recent works approach this problem under a game theoretic lens, whereby utility functions are assigned to each of the agents with the hope that their local optimization approximates the centralized optimal solution. Yet, formal guarantees on the resulting performance cannot be obtained for broad classes of problems without compromising on their accuracy. In this work, we address this concern relative to the well-studied problem of resource allocation with nondecreasing concave welfare functions. We show that optimally designed local utilities achieve an approximation ratio (price of anarchy) of $1 - c/e$, where c is the function’s curvature and e is Euler’s constant. The upshot of our contributions is the design of approximation algorithms that are distributed and efficient, and whose performance matches that of the best existing polynomial-time (and centralized) schemes.

KEYWORDS

distributed submodular maximization, approximation ratio, price of anarchy, game theory, resource allocation

1 INTRODUCTION

The study of distributed control in multiagent systems has gained popularity over the past few decades as it has become apparent that the behaviour of local decision makers impacts the performance of many social and technological systems. Consider the typical example of selfish drivers on a road network. Counterintuitively, if all drivers make route selections that minimize their own travel times, the average time each driver spends on the road can be much higher than optimal [10]. As an alternative example, DARPA’s Blackjack program aims to launch satellite constellations with a high degree of mission-level autonomy into low Earth orbit [14]. The key objectives of the Blackjack program include developing the on-orbit, distributed decision making capabilities within these satellite networks, as agent coordination cannot rely upon the unreliable and high latency communications from ground control.

In either of the scenarios described above, the system would perform most efficiently if a central coordinator could compute and relay the optimal decisions to each of the agents. However, in the systems we have discussed, coordination by means of a central authority is either impractical – due to latencies and bandwidth limitations in communications, scalability and security requirements, etc. – or even impossible (e.g., dictating what route each driver must follow is not presently possible). Thus, in these systems, decision making *must* be distributed. The inevitable loss in performance

when coordination is distributed – often referred to as the “tragedy of the commons” in economics and environmental sciences [17, 22] – is well-documented in many scenarios [1, 10, 16]. Evidently, the design of algorithms that mitigate the losses in system performance stemming from the distribution of decision making is critical to the implementation of the multiagent systems described.

A fruitful paradigm for the design of distributed multiagent coordination algorithms – termed the *game theoretic approach* [24, 31] – involves modelling the agents as players in a game and assigning them utility functions that maximize the efficiency of the game’s equilibria. After agents’ utilities are coupled with learning dynamics capable of driving the system to an equilibrium, an efficient distributed coordination algorithm emerges. This approach has been utilized in a variety of relevant contexts, including collaborative sensing in distributed camera networks [12], the distributed control of smart grid nodes [30], autonomous vehicle-target assignment [2] and optimal taxation on road-traffic networks [26]. A significant advantage of such an approach is that the design of the agents’ learning dynamics and of the underlying utility structure can be decoupled. As efficient distributed learning dynamics that drive the agents to an equilibrium of the game are already known (see, e.g. [18]), we focus our attention on the design of agents’ utility functions in order to maximize the efficiency of the equilibria.

The most commonly studied metric in the literature on utility design is the *price of anarchy* [21], which is defined as the worst case ratio between the performance at an equilibrium and the best achievable system performance. Note that a price of anarchy guarantee obtained for a set of utility functions translates directly to an approximation ratio of the final distributed algorithm. The majority of the literature focuses primarily on characterizing the price of anarchy for a given set of player utility functions [23, 24, 34], whereas fewer works design player utilities in order to optimize the price of anarchy [8, 15, 27]. While several works provide tight bounds on the approximation ratio of polynomial-time centralized algorithms for the class of problems we consider (see, e.g., [4, 13, 32]), there is currently no result in the literature that establishes comparable bounds on the best achievable price of anarchy, aside from the general bound put forward in Vetta [34] that is provably inexact.

1.1 Model

In this paper, we consider a class of resource allocation problems with a set of agents $N = \{1, \dots, n\}$ and a set of resources \mathcal{R} . Each resource $r \in \mathcal{R}$ has a corresponding welfare function $W_r : \mathbb{N} \rightarrow \mathbb{R}$. Each agent $i \in N$ must select an action a_i from a corresponding set of actions $\mathcal{A}_i \subseteq 2^{\mathcal{R}}$. The system performance under an allocation of agents $a = (a_1, \dots, a_n) \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ is measured by a function $W : \mathcal{A} \rightarrow \mathbb{R}_{>0}$. The goal is to find an allocation $a^{\text{opt}} \in \mathcal{A}$

that maximizes the function

$$W(a) := \sum_{r \in \cup_i a_i} W_r(|a|_r), \quad (1)$$

where $|a|_r = |\{i \in N \text{ s.t. } r \in a_i\}|$ denotes the number of agents selecting the resource r in allocation a . In this work, we consider nonnegative, nondecreasing concave welfare functions, i.e., functions that satisfy the following properties: (i) $W_r(x)$ is nondecreasing and concave for $x \geq 0$; and, (ii) $W_r(0) = 0$ and $W_r(x) > 0$ for all $x \geq 1$. This setup has been thoroughly studied in the submodular maximization and game theoretic literature (see, e.g., [2–4, 15, 32]) as demonstrated by the following two examples:

Example 1 (General covering problems). Consider the general covering problem [15], which is a generalization of the max-n-cover problem [13, 19]. In this setting, we are given a set of elements E and n collections S_1, \dots, S_n of subsets of E , i.e., $S_i \subseteq 2^E$ for all $i = 1, \dots, n$. Each element $e \in E$ has weight $w_e \geq 0$. The objective is to choose one subset s_i from each collection S_i such that the union $\cup_i s_i$ has maximum total weight, i.e., $\sum_{e \in \cup_i s_i} w_e$ is maximized. We observe that this problem corresponds to a resource allocation problem where each agent $i \in N$ has action set $\mathcal{A}_i \subseteq 2^E$, the action a_i of each agent $i \in N$ corresponds to the subset s_i , and the welfare functions are $W_e(x) = w_e$ for all $e \in E$.

Example 2 (Vehicle-target assignment problem). Consider the *vehicle-target assignment problem*, first introduced in Murphey [25], and studied by, e.g., Arslan et al. [2] and Barman et al. [3]. In this setting, we are given a set of n vehicles N and a set of targets \mathcal{T} , where each target $t \in \mathcal{T}$ has an associated value $v_t > 0$. Each vehicle $i \in N$ has a set of feasible target assignments $\mathcal{A}_i \subseteq 2^{\mathcal{T}}$. Given that a vehicle $i \in N$ is assigned to target $t \in \mathcal{T}$, the probability that t is destroyed by i is $p_t \in (0, 1]$. The objective is to compute a joint assignment of vehicles $a \in \prod_i \mathcal{A}_i$ that maximizes the expected value of targets destroyed, which is measured as

$$W(a) = \sum_{t \in \mathcal{T}} v_t \cdot (1 - (1 - p_t)^{|a|_t}), \quad (2)$$

where $1 - (1 - p_t)^x$ is the probability that target t is destroyed when x vehicles are assigned to it. Observe that the vehicle-target assignment problem is a resource allocation problem with nonnegative, nondecreasing concave welfare functions where the agents are the vehicles, the resources are the targets, and the welfare function on each resource $t \in \mathcal{T}$ is $W_t(x) = v_t \cdot (1 - (1 - p_t)^x)$.¹

The focus of this work is on computing near-optimal distributed solutions within the class of resource allocation problems described above using the game theoretic approach. To model this particular class of problems, we adopt the framework of *resource allocation games*. A resource allocation game $G = (N, \mathcal{R}, \mathcal{A}, \{F_r\}_{r \in \mathcal{R}})$ consists of a player set N where each player $i \in N$ evaluates the allocation $a \in \mathcal{A}$ using a utility function

$$U_i(a) := \sum_{r \in a_i} F_r(|a|_r). \quad (3)$$

¹Observe that a vehicle-target assignment problem with $p_t = p = 1.0$ for all targets $t \in \mathcal{T}$ is equivalent to a general covering problem where the chosen subsets correspond with the vehicle assignments and the element weights are equal to the target values. Nonetheless, we retain Example 1 as it is thoroughly connected to the literature.

where $F_r : \mathbb{N} \rightarrow \mathbb{R}$ defines the utility a player receives at resource r as a function of the total number of agents selecting r in allocation a . We refer to the functions $\{F_r\}_{r \in \mathcal{R}}$ as the *local utility functions* of the game. For a given set of welfare functions \mathcal{W} , it is convenient to define a utility mapping $\mathcal{F} : \mathcal{W} \rightarrow \mathbb{R}^N$, where it is understood that a resource r with welfare function $W_r \in \mathcal{W}$ is assigned the local utility function $\mathcal{F}(W_r)$.

In the forthcoming analysis, we consider the solution concept of pure Nash equilibrium, which is defined as any allocation $a^{\text{ne}} \in \mathcal{A}$ such that

$$U_i(a^{\text{ne}}) \geq U_i(a_i, a_{-i}^{\text{ne}}), \quad \forall a_i \in \mathcal{A}_i, \quad \forall i \in N, \quad (4)$$

where $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n)$. For a given game G , let $\text{NE}(G)$ denote the set of all allocations $a \in \mathcal{A}$ that satisfy the Nash condition in Equation (4). We define the *price of anarchy* of a game G as²

$$\text{PoA}(G) := \frac{\min_{a \in \text{NE}(G)} W(a)}{\max_{a \in \mathcal{A}} W(a)} \leq 1. \quad (5)$$

For a given game G , the price of anarchy is the ratio between the system-wide performance of the worst performing pure Nash equilibrium and the optimal allocation. The price of anarchy as defined here also applies to the efficiency of the game’s coarse-correlated equilibria [8, 29], for which many efficient algorithms exist (see, e.g., [18]). We extend the definition of price of anarchy to a given set of games \mathcal{G} , which may contain infinitely many instances, as $\text{PoA}(\mathcal{G}) := \inf_{G \in \mathcal{G}} \text{PoA}(G) \leq 1$. It is important to note that a higher price of anarchy corresponds to an overall improvement in the performance of all pure Nash equilibria, and that $\text{PoA}(\mathcal{G}) = 1$ implies that all pure Nash equilibria in all games $G \in \mathcal{G}$ are optimal. For a given utility mechanism \mathcal{F} , we use the terminology “the set of games \mathcal{G} induced by the set of welfare functions \mathcal{W} ” to refer to the set of all games with $W_r \in \mathcal{W}$ and $F_r = \mathcal{F}(W_r)$ for all $r \in \mathcal{R}$. Given a set \mathcal{W} , our aim is to develop an efficient technique for computing a utility mechanism \mathcal{F}^{opt} that maximizes the price of anarchy in the corresponding set of games \mathcal{G} induced by \mathcal{W} , i.e., we wish to solve

$$\mathcal{F}^{\text{opt}} \in \arg \max_{\mathcal{F}} \text{PoA}(\mathcal{G}). \quad (6)$$

1.2 Results and Discussion

Our main result is an efficient technique for computing a utility mechanism that guarantees a price of anarchy of $1 - c/e$ in all resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature c .

Definition 1 (Curvature [11]). The curvature of a nondecreasing concave function $W : \mathbb{N} \rightarrow \mathbb{R}$ is

$$c = 1 - \frac{W(n) - W(n-1)}{W(1)}. \quad (7)$$

In the literature on submodular maximization, the curvature is commonly used to compactly parameterize broad classes of functions. The notion of curvature we consider was originally defined by Conforti et al. [11] in the context of general nondecreasing submodular set functions. In our specific setup, this reduces to the expression in Definition 1. Observe that all nondecreasing concave functions

²Note that the price of anarchy is well-defined for resource allocation games, since these games possess a potential function and, thus, at least one pure Nash equilibrium.

have curvature $c \in [0, 1]$. Thus, $c = 1$ can be considered in scenarios where the maximum curvature among functions in the set \mathcal{W} is not known.

Theorem 1 (Informal). *Let \mathcal{G} denote the set of all resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature c . An optimal utility mechanism achieves $\text{PoA}(\mathcal{G}) = 1 - c/e$ and can be computed efficiently.*

A significant consequence of the main result is a universal guarantee that the best achievable price of anarchy is always greater than $1 - 1/e \approx 63.2\%$ for resource allocation games with nonnegative, nondecreasing concave welfare functions. Note that since $1 - 1/e$ is the optimal price of anarchy in general covering games (see, e.g., Example 1), it cannot be further improved without more information about the underlying set of welfare functions. Our guarantee improves to $1 - c/e$ if the curvature c of the underlying set of welfare functions is known.

Observe that the result in Theorem 1 also implies that one can efficiently compute a “universal” utility mechanism, in that it would guarantee a price of anarchy greater than or equal to $1 - 1/e$ with respect to any game with nonnegative, nondecreasing concave welfare functions. This follows from the observation that $c \leq 1$ always holds. Of course, if more information is available about the underlying set of welfare functions (e.g., the maximum curvature), then this lower bound can be improved. In the case where the entire set of welfare functions \mathcal{W} is known *a priori* and $|\mathcal{W}|$ is “small enough”, then the optimal utility mechanism can be computed using existing methodologies (see, e.g., [8]).³ Consider the sets represented in Figure 1. From our reasoning, it holds that as the size of the set of welfare functions considered is reduced, the prices of anarchy of the corresponding optimal utility mechanisms increase. The set of games induced by welfares in the green ellipse, for example, coincides with the vehicle-target assignment problem, as described in Example 2, where $p_t = p \in [0, 1]$ for all $t \in \mathcal{T}$. Note that the welfare function W_t of each target $t \in \mathcal{T}$ in this problem is nonnegative, nondecreasing concave (i.e., the green ellipse is a subset of the dotted red box). Thus, we can immediately observe that the best achievable price of anarchy in the corresponding resource allocation game G satisfies

$$\text{PoA}(G) \geq 1 - \frac{1}{e},$$

which is achieved by the universal utility mechanism from Theorem 1. Since there is only a single welfare function in this setting (ignoring uniform scalings) the optimal utility mechanism can be computed for a modest number of agents, as aforementioned.

In Figure 2, we plot the price of anarchy corresponding to the optimal utility mechanism within this setting (labelled “Optimal”), the price of anarchy achieved by the universal utility mechanism (labelled “Universal”) and the $1 - 1/e$ lower bound from Theorem 1 (labelled “Lower bound”). As expected, the optimal utility mechanism corresponds with the best price of anarchy as it was designed specifically for the underlying welfare function. However, knowledge of the set of welfare functions corresponds with only a small

³In this case, the optimal utility mechanism can be found as the solution of $|\mathcal{W}|$ linear programs with number of constraints that is quadratic in the maximum number of agents n , and $n + 1$ decision variables. For this reason, the optimal utility mechanism can only be computed for modest values of $|\mathcal{W}|$ and n .

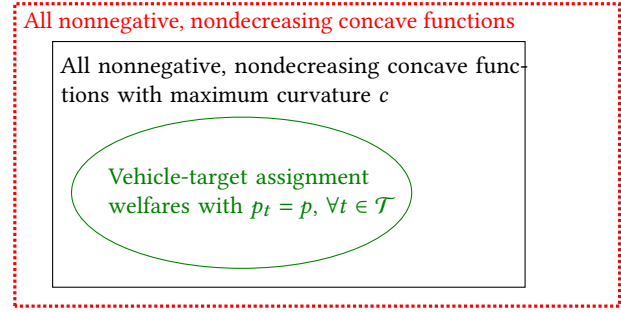


Figure 1: The set of games induced by the set of all nonnegative, nondecreasing concave functions contains the set of all nonnegative, nondecreasing concave functions with maximum curvature c , which in turn contains the set of all vehicle-target assignment problems with $p_t = p$.




Figure 2: The price of anarchy of the universal utility mechanism obtained in this work and the optimal utility mechanism in the vehicle-target assignment problems with $p_t \in [0, p]$ for all $t \in \mathcal{T}$. Note that this utility mechanism is designed for the set of all nonnegative, nondecreasing concave welfare functions but its price of anarchy is close to the best achievable within this particular setting.

increase in the price of anarchy; the price of anarchy achieved by the universal utility mechanism is surprisingly close to the best achievable by any mechanism for all values of $p \in [0, 1]$. Note that the universal utility mechanism is only guaranteed to achieve a price of anarchy of $1 - 1/e$.

Consider once again the sets represented in Figure 1. The above example suggests that the utility mechanism designed to maximize the price of anarchy in the set of games induced by all welfare functions in the dotted, red box may achieve price of anarchy close to the optimal within the sets of games induced by any subset of the dotted red box, as we have observed that this holds for the set of games induced by welfare functions in the green ellipse. While we do not provide formal proofs for these observations, they provide further motivation for deriving efficient techniques for computing utility mechanisms that maximize the price of anarchy with respect to broad classes of welfare functions.

1.3 Related Works

Submodular resource allocation problems have been the focus of a significant research effort for many years, particularly in the optimization community. Since the computation of an optimal allocation in such problems is \mathcal{NP} -hard in general, many researchers have focused on providing approximation guarantees for polynomial-time algorithms. For example, approximate solutions to max- n -cover problems were studied by Feige [13] and Hochbaum [19] almost 25 years ago. In the latter manuscript, the greedy algorithm is shown to have an approximation ratio of $1 - 1/e$. Recently, Sviridenko et al. [32] proposed a polynomial-time algorithm for computing approximate solutions that perform within $1 - c/e$ of the optimal for the class of resource allocation problems with nonnegative, nondecreasing submodular welfare functions with curvature c . Barman et al. [4] provide a polynomial-time algorithm that returns allocations with a $1 - k^k e^{-k} / (k!)$ approximation ratio in resource allocation problems with welfare functions $W_r(x) = \min\{x, k\}$ for $k \in \mathbb{N}_{\geq 1}$, for all $r \in \mathcal{R}$. In their respective works, all three of the approximation ratios provided above are also shown to be the best achievable, i.e., it is shown that there exist no other polynomial-time algorithm capable of always computing an approximate solution that is closer to the optimal unless $\mathcal{P} = \mathcal{NP}$. In this work, we obtain price of anarchy guarantees in resource allocation games that match these approximation ratios from submodular maximization.

Although utility mechanisms have been studied in resource allocation games, the majority of results have focused on deriving price of anarchy bounds for given utility structures (e.g., marginal contribution, equal shares, etc.) [23, 24]. In this respect, Vetta [34] proves that there always exist player utility functions that guarantee a price of anarchy larger than 50% within a more general class of games than those we consider here.⁴ The notion of utility mechanisms that *maximize* – or otherwise improve – the price of anarchy was introduced in Christodoulou et al. [9]. This approach has been applied to many distributed optimization problems, including machine scheduling [6, 20], selfish routing [5, 7] and auction mechanism design [28, 33]. A prominent example in this line of research is Gairing [15], who proves that the best achievable price of anarchy in covering games is $1 - 1/e$ and derived an optimal utility mechanism. We provide an efficient technique for computing a utility mechanism that achieves a price of anarchy larger than $1 - 1/e \approx 63.2\%$ in all resource allocation games with nonnegative, nondecreasing concave welfare functions, which effectively generalizes the result in [15] and significantly improves upon the bound provided in [34] by exploiting the structure on W (see Equation (1)).

More recently, Chandan et al. [8] proposed a linear programming based approach for computing a utility mechanism that maximizes the price of anarchy for a given set of resource allocation games. Unfortunately, their approach does not provide *a priori* guarantees on the price of anarchy achieved. In other words, a concrete lower bound on the best achievable price of anarchy using their approach cannot be obtained without first solving the linear programs corresponding to the underlying set of welfare functions. This can only be accomplished with this linear programming based approach

⁴In the class of *valid-utility games*, the system objective $W : \mathcal{A} \rightarrow \mathbb{R}$ is a nondecreasing submodular set function over the agents' actions and is not necessarily separable over a set of resources; much more general than the class of resource allocation games with nonnegative, nondecreasing concave welfare functions.

for a finite number of welfare functions and a modest number of players. Our goal in this work is to derive guarantees on the best achievable price of anarchy for games induced by the set of all welfare functions satisfying a few specific properties (e.g., concavity, maximum curvature). In the next section, we provide an explicit expression for a utility mechanism that is guaranteed to have price of anarchy greater than or equal to $1 - c/e$ for all resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature c . Since the linear programming based approach returns a utility mechanism that maximizes the price of anarchy, the claim in Theorem 1 extends to the solutions of the linear program.

1.4 Organization

The remainder of the paper is structured as follows: Section 2 presents the proof of the main result and an extension result for more specific sets of welfare functions. Section 3 showcases our simulation example and accompanying discussion. Section 4 concludes the manuscript and provides a brief discussion on potential future directions. All proofs omitted from the manuscript are provided in the appendix, for ease of exposition.

2 MAIN RESULT AND EXTENSIONS

In this section, we prove the claim in Theorem 1 by constructing a utility mechanism that achieves the best achievable price of anarchy of $1 - c/e$ with respect to the set of all nonnegative, nondecreasing concave welfare functions with maximum curvature $c \in [0, 1]$. In scenarios where a more specific set of welfare functions is considered, we outline how the techniques used to prove Theorem 1 can be generalized to derive tighter *a priori* bounds on the best achievable price of anarchy.

Before presenting the proof of Theorem 1, we provide an informal outline of the three steps underpinning the result. These steps correspond with the three parts of the formal proof, but are listed in a different order for sake of clarity. For the reader's convenience, we include the part of the proof that corresponds with each of the steps in our informal outline. The proof is summarized as follows:

–*Step #1:* We demonstrate that any concave welfare function can be decomposed as a linear combination with nonnegative coefficients of a specialized set of basis functions. [Section 2.1, Part ii)]

–*Step #2:* We derive optimal basis utility functions for each of the basis functions in the specialized set. [Section 2.1, Part i)]

–*Step #3:* We construct local utility functions as linear combinations over the optimal basis utility functions from Step 2 with the nonnegative coefficients derived in Step 1. Finally, we demonstrate that this tractable approach for constructing resource utility functions provides near optimal efficiency guarantees. [Section 2.1, Part iii)]

2.1 Proof of Theorem 1

Here we consider the class of games induced by the set of all concave welfare functions with maximum curvature $c \in [0, 1]$. The proof of Theorem 1 proceeds in the following three parts:

- i) Given a value $c \in [0, 1]$, we derive explicit expressions for the local utility functions that maximize the price of anarchy relative to a restricted class of nonnegative, nondecreasing concave welfare functions with curvature c . Among the optimal price of anarchy values obtained for the functions in this restricted class, the lowest is equal to $1 - c/e$;
- ii) We show that any nonnegative, nondecreasing concave welfare function W with curvature less than or equal to c can be represented as a linear combination with explicitly defined nonnegative coefficients over this restricted class; and,
- iii) We demonstrate that using the local utility functions computed as a linear combination over the optimal local utility functions from i) with the nonnegative coefficients from ii) guarantees that $\text{PoA}(\mathcal{G}) = 1 - c/e$ within the set of resource allocation games \mathcal{G} induced by all nonnegative, nondecreasing concave welfare functions with maximum curvature c .

The above parts successfully prove Theorem 1 as we argue here. Note that, by part i), the lowest optimal price of anarchy among welfare functions in the restricted class considered is equal to $1 - c/e$, for given curvature $c \in [0, 1]$. By part iii), this implies that all resource allocation games induced by nonnegative, nondecreasing concave welfare functions with maximum curvature c have optimal price of anarchy equal to $1 - c/e$. This is because, by part ii), any such welfare function can be represented as a nonnegative linear combination over the restricted class of welfare functions we consider. Since the best achievable price of anarchy for at least one of the functions in the restricted class is also $1 - c/e$, one cannot further improve the price of anarchy within the set of games considered. In addition, parts i)–iii) combine to prove that a corresponding utility mechanism that maximizes the price of anarchy entails computing nonnegative linear combinations over a class of functions with explicit expressions. Thus, the computation of optimal local utility functions is polynomial in the number of players.

Part i). In this part of the proof, we provide explicit expressions for local utility functions that maximize the price of anarchy with respect to a restricted set of welfare functions, as well as the corresponding optimal price of anarchy. To that end, given parameters $\alpha \in [0, 1]$ and $\beta \in \mathbb{N}_{\geq 1}$, we define the (α, β) -coverage function as

$$V_{\beta}^{\alpha}(x) := (1 - \alpha) \cdot x + \alpha \cdot \min\{x, \beta\}. \quad (8)$$

It is straightforward to verify that every (α, β) -coverage function is nonnegative, nondecreasing concave. In the lemma below, we derive a local utility function that maximizes the price of anarchy of the set of resource allocation games induced by any given (α, β) -coverage function. We use this result to derive the optimal utility functions for a broad range of local welfare functions in Part iii).

Lemma 1. Consider the set of resource allocation games \mathcal{G} induced by the (α, β) -coverage function

$$V_{\beta}^{\alpha}(x) = (1 - \alpha) \cdot x + \alpha \cdot \min\{x, \beta\},$$

where $\alpha \in [0, 1]$ and $\beta \in \mathbb{N}_{\geq 1}$. Let $\rho = (1 - \alpha \cdot \beta^{\beta} e^{-\beta} / (\beta!))^{-1}$, and define F_{β}^{α} as in the following recursion: $F_{\beta}^{\alpha}(1) := W(1)$,

$$F_{\beta}^{\alpha}(x+1) := \max \left\{ \frac{1}{\beta} [x F_{\beta}^{\alpha}(x) - V_{\beta}^{\alpha}(x) \rho] + 1, 1 - \alpha \right\}, \forall x = 1, \dots, n-1. \quad (9)$$

Then, the local utility function F_{β}^{α} maximizes the price of anarchy and the corresponding price of anarchy is $\text{PoA}(\mathcal{G}) = 1/\rho$.

PROOF. The proof is presented in Appendix A.1. \square

According to the result in Lemma 1, the maximum achievable price of anarchy in resource allocation games induced by a (α, β) -coverage function with $\alpha = 1$ and $\beta \geq 1$ is $1 - \beta^{\beta} e^{-\beta} / (\beta!)$. Surprisingly, Barman et al. [4] show that the optimal approximation ratio of any polynomial-time algorithm for the same class of resource allocation problems is also $1 - \beta^{\beta} e^{-\beta} / (\beta!)$. Similarly, the optimal price of anarchy for the (α, β) -coverage function with $\alpha \in [0, 1]$ and $\beta = 1$ is $1 - \alpha/e$, which matches the best achievable approximation ratio of any polynomial-time algorithm for this problem setting [32].

Part ii). In the next result, we show that any nonnegative, nondecreasing concave welfare function with maximum curvature $c \in [0, 1]$ can be represented as a nonnegative linear combination over the set of (c, k) -coverage functions with $k = 1, \dots, n$.

Lemma 2. Let $W : \mathbb{N} \rightarrow \mathbb{R}$ denote a nonnegative, nondecreasing concave function with curvature less than or equal to $c \in [0, 1]$. Then, the nonnegative coefficients η_1, \dots, η_n satisfy

$$W(x) = \sum_{k=1}^n \eta_k \cdot V_k^c(x), \quad \forall x = 0, 1, \dots, n, \quad (10)$$

where $\eta_1 := [2W(1) - W(2)]/c$, $\eta_k := [2W(k) - W(k-1) - W(k+1)]/c$, for $k = 2, \dots, n-1$, and $\eta_n := W(1) - \sum_{k=1}^{n-1} \eta_k$.

PROOF. The proof is presented in Appendix A.3. \square

Part iii). We begin by describing a utility mechanism parameterized by the maximum curvature and maximum number of players. Let \mathcal{G} denote the set of resource allocation games induced by all nonnegative, nondecreasing concave functions with maximum curvature $c \in [0, 1]$ with a maximum of n players. Consider any resource allocation game $G \in \mathcal{G}$ and assign the following local utility function to each $r \in \mathcal{R}$:

$$F_r(x) = \sum_{k=1}^n \eta_k \cdot F_k^c(x), \quad \forall x = 1, \dots, n,$$

where $\eta_1 := [2W_r(1) - W_r(2)]/c$, $\eta_k := [2W_r(k) - W_r(k-1) - W_r(k+1)]/c$, for $k = 2, \dots, n-1$, and $\eta_n := W_r(1) - \sum_{k=1}^{n-1} \eta_k$, $W_r : \mathbb{N} \rightarrow \mathbb{R}$ is the welfare function on the resource r and each function $F_k^c : \mathbb{N} \rightarrow \mathbb{R}$, $k = 1, \dots, n$, is the optimal local utility function for $V_k^c(x)$ defined recursively in Lemma 1. In this part, we show that $\text{PoA}(G) \geq 1 - c/e$ holds for this utility mechanism.

Given maximum curvature $c \in [0, 1]$, Lemma 1 proves that among the (c, k) -coverage functions with $k = 1, \dots, n$, the $(c, 1)$ -coverage function has best achievable price of anarchy $1 - c/e$ which is strictly lower than the best achievable price of anarchy for any (c, k) -coverage function with $k > 1$. This implies that the best achievable price of anarchy must satisfy $\text{PoA}(\mathcal{G}) \leq 1 - c/e$, since any game G in the set of resource allocation games induced by the $(c, 1)$ -coverage function must also be in the set \mathcal{G} , i.e., $G \in \mathcal{G}$, and there is at least one such game with $\text{PoA}(G) = 1 - c/e$. We now

show that $\text{PoA}(\mathcal{G}) \geq 1 - c/e$ also holds. Recall from Lemma 2 that the nonnegative coefficients η_1, \dots, η_n defined above satisfy

$$W_r(x) = \sum_{k=1}^n \eta_k \cdot V_k^c(x) \quad \forall x = 0, 1, \dots, n.$$

It must then hold that, for any $r \in \mathcal{R}$, $(F_r, (1 - c/e)^{-1})$ is a feasible point in the linear program in Equation (12) (see Appendix A) for any n and the corresponding W_r . Observe that each constraint in the linear program must be satisfied since, by Lemma 2, it can be represented as a nonnegative linear combination of the constraints in the n linear programs for V_k^c and $(F_k^c, (1 - c/e)^{-1})$, $k = 1, \dots, n$, i.e., for all $r \in \mathcal{R}$ and all $(x, y, z) \in \mathcal{I}(n)$ it must hold that

$$\begin{aligned} (1 - c/e)^{-1} W_r(x) &\geq \sum_{k=1}^n \eta_k \cdot \left[1 - c \cdot \frac{k^k e^{-k}}{k!} \right] V_k^c(x) \\ &\geq \sum_{k=1}^n \eta_k \cdot [V_k^c(y) + (x - z)F_k^c(x) - (y - z)F_k^c(x + 1)] \\ &= W_r(y) + (x - z)F_r(x) - (y - z)F_r(x + 1), \end{aligned}$$

where the first inequality holds because $1 - c/e \leq 1 - c \cdot k^k e^{-k} / (k!)$ for all $k \geq 1$ and since $W_r, V_k^c(x)$, $k = 1, \dots, n$, and the coefficients η_1, \dots, η_n are nonnegative, and the second inequality holds by the result in Lemma 1.

2.2 Specialized sets of welfare functions

In the previous subsection, we used a series of arguments to prove the bound on the price of anarchy in Theorem 1. Informally, we considered a specified set of candidate welfare functions. For this set of candidate welfare functions, we derived a corresponding set of local utility functions that maximize the price of anarchy. Finally, we showed that the best achievable price of anarchy for these candidates is automatically a lower bound on the best achievable price of anarchy across a much broader set of welfare functions. A set of candidate welfare functions must be chosen for two reasons: (i) an optimal local utility function and its corresponding optimal price of anarchy can be obtained in advance for each of the candidate welfare functions; and, more importantly, (ii) any function within the set of welfare functions of interest can be expressed as a nonnegative linear combination over the set of candidate welfare functions, thus inheriting the same optimal price of anarchy. Clearly the choice of candidate functions is important, as the *a priori* guarantees on the price of anarchy is characterized by the best achievable price of anarchy corresponding to each candidate.

As our next result, we outline a mechanism for obtaining a set of candidate functions for a given set of welfare functions \mathcal{W} such that any function $W \in \mathcal{W}$ can be expressed as a nonnegative linear combination over the candidate functions. This generalizes the approach taken in the previous subsection to sets of resource allocation games for which more is known about the welfare functions than concavity and maximum curvature $c \in [0, 1]$.

Corollary 1. *Let \mathcal{W} denote a set of nonnegative, nondecreasing concave welfare functions and n be the maximum number of agents. Let W^{ub} and W^{lb} be two nonnegative, nondecreasing concave functions that satisfy the following for all $W \in \mathcal{W}$: (i) $W^{lb}(x + 1) - W^{lb}(x) \leq [W(x + 1) - W(x)]/W(1) \leq W^{ub}(x + 1) - W^{ub}(x)$, for all $x =$*

1, \dots, n - 1; and, (ii) $[W(x + 1) - 2W(x) + W(x - 1)]/W(1) \leq W^{ub}(x + 1) - 2W^{ub}(x) + W^{ub}(x - 1) \leq W^{lb}(x + 1) - 2W^{lb}(x) + W^{lb}(x - 1)$, for all $x = 2, \dots, n - 1$. Finally, define the candidate functions $W^{(k)}$, $k = 1, \dots, n$, as follows:

$$W^{(k)}(x) = \begin{cases} W^{ub}(x) & \text{if } 1 \leq x \leq k, \\ W^{ub}(k) + W^{lb}(x) - W^{lb}(k) & \text{if } x > k. \end{cases} \quad (11)$$

Then, for any welfare function $W \in \mathcal{W}$, there exist nonnegative coefficients η_1, \dots, η_n that satisfy

$$W(x) = \sum_{k=1}^n \eta_k \cdot W^{(k)}(x), \quad \forall x = 0, 1, \dots, n.$$

PROOF. The proof is in Appendix B. \square

We highlight several important implications of the result in Corollary 1 in the following discussion:

(i) We showed in Part iii) of the previous subsection that any set of resource allocation games \mathcal{G} induced by nonnegative linear combinations over a set of candidate functions $W^{(1)}, \dots, W^{(n)}$ automatically inherits the optimal price of anarchy guarantees of the candidates, i.e., there exist local utility functions such that $\text{PoA}(\mathcal{G})$ is greater than or equal to the lowest optimal price of anarchy among the candidates. Thus, by simply precomputing the optimal local utility functions $F^{(1)}, \dots, F^{(n)}$ and price of anarchy bounds corresponding to the candidate functions, one obtains a lower bound on the best achievable price of anarchy in the set of games considered. This can be done, for example, using the linear programming based methodology proposed in [8].

(ii) If the candidate function with lowest corresponding optimal price of anarchy happens to be a member of the underlying set \mathcal{W} , then we can also say that this lower bound is the best achievable price of anarchy. Furthermore, an optimal utility mechanism then consists of computing nonnegative linear combination over the precomputed functions $F^{(1)}, \dots, F^{(n)}$.

(iii) The complexity of computing the local utility functions that achieve the lower bound on $\text{PoA}(\mathcal{G})$ is polynomial in the number of players. This follows from observing that the functions $F^{(1)}, \dots, F^{(n)}$ can be precomputed and there is a closed-form expression for the nonnegative coefficients η_k , $k = 1, \dots, n$, given a welfare function $W \in \mathcal{W}$ (see, e.g., the proof of Corollary 1).

3 SIMULATION RESULTS

In this section, we provide an in-depth simulation example in which we compare the equilibrium performance corresponding to the universal utility mechanism we derive in the previous section for $c = 1$ against two well-studied utility structures from the literature: the *identical interest utility* and the *equal shares utility mechanism*. The identical interest utility precisely aligns the players' utilities to the system objective, i.e., $U_i(a) = W(a)$ for all $i \in N$. Observe that under this utility, if $U_i(a_i, a_{-i}) > U_i(a'_i, a_{-i})$ for a player $i \in N$, then it must hold that $W(a_i, a_{-i}) > W(a'_i, a_{-i})$. As its name suggests, the equal shares utility mechanism distributes the welfare obtained on each resource among the players selecting that resource which corresponds with local utility functions of the form $F_r^{\text{es}}(x) = W_r(x)/x$




Figure 3: Box plots depicting the equilibrium efficiency measured across $T = 10^3$ instances for the universal utility, identical interest utility and equal shares utility mechanisms in the vehicle-target assignment problem with $p_t = p$ for all $t \in \mathcal{T}$ and $p \in \{0.5, 0.6, 0.7\}$. Note that among the three utility mechanisms studied, the price of anarchy is highest for the universal utility mechanism.

for all $r \in \mathcal{R}$. At first glance, one might expect that one of these two utilities would be best, e.g., the identical interest utility exposes the players to the actual system objective. However, in terms of the worst-case equilibrium efficiency, our simulation provides concrete evidence that the universal utility mechanism performs better.

Consider a vehicle-target assignment problem with $n = 10$ vehicles and $|\mathcal{T}| = n + 1$ targets, where $\mathcal{T} = \{t_1, \dots, t_{n+1}\}$. We purposely choose a small number of vehicles (i.e., $n = 10$) in order to allow for explicit computation of the optimal allocation and, therefore, of the corresponding price of anarchy. Each vehicle $i \in N$ has two singleton target assignments chosen randomly from a uniform distribution over the $n + 1$ targets, i.e., $\mathcal{A}_i = \{\{t_j\}, \{t_k\}\}$ where $j, k \sim \mathcal{U}\{1, n + 1\}$. Each target $t \in \mathcal{T}$ has welfare function $W_t(x) = v_t \cdot (1 - (1 - p)^x)$ where v_t is drawn from a uniform distribution over the interval $[0, 1]$ and $p \in [0, 1]$ is a given parameter.

Within the scenario described above, we model agent decision making as best response dynamics over $T = 100$ iterations. More specifically, the agents best respond in a round robin fashion to the actions of the others, i.e., at each time step $t \in \{1, \dots, T\}$, the agent $i = t \bmod n$ selects an action $a_i^t \in \mathcal{A}_i$ such that $U_i(a_i^t, a_{-i}^{t-1}) = \max_{a_i \in \mathcal{A}_i} U_i(a_i, a_{-i}^{t-1})$, and then $a^t = (a_i^t, a_{-i}^{t-1})$. As the agents settled to a pure Nash equilibrium within 20 iterations in all the instances we generated, repeating over $T = 100$ iterations is justified. We ran our simulations for the three utility structures described (i.e., universal utility mechanism, identical interest utility and equal shares utility mechanism) over 10^3 randomly generated instances for $p \in \{0.5, 0.6, 0.7\}$, as described above. The price of anarchy data was obtained by dividing the welfare at equilibrium by the best achievable welfare computed by exhaustive search. The box plots in Figure 3 display statistics on the price of anarchy values we obtain

in our simulations. These box plots are to be interpreted as follows: (i) the top and bottom of the boxes correspond to the 75-th and 25-th percentiles of the price of anarchy, respectively; (ii) the top and bottom “whiskers” show the maximum and minimum price of anarchy, respectively; and, (iii) each of the boxes is bisected by the median value of the corresponding prices of anarchy.

Observe that for all three values of p considered, the minimum price of anarchy across the 10^3 randomly generated instances is highest for the universal utility mechanism, as expected. However, for all three utility functions considered, the maximum and 75-th percentile of the price of anarchy data collected is always at 1, i.e., the best response dynamics settled on an optimal allocation for at least 25% of the randomly generated instances. In fact, all of the other statistics on the price of anarchy are skewed away from the minimum, suggesting that the worst-case instances are quite rare. Furthermore, although the minimum price of anarchy for the identical interest utility is lowest for $p \in \{0.5, 0.6, 0.7\}$, the identical interest utility also has the highest median and 25-th percentile price of anarchy values among the three utilities considered. These observations suggest that – as one might expect – the price of anarchy is not representative of the average equilibrium efficiency, and that the identical interest utility could perform better than the universal utility mechanism in this respect. The design of utility functions that maximize the expected equilibrium efficiency could be a fruitful direction for future work.

4 CONCLUSIONS AND OPEN QUESTIONS

In this work, we consider the game theoretic approach to the design of distributed algorithms for resource allocation problems with nonnegative, nondecreasing concave welfare functions. Our main result is that there exist utility mechanisms that achieve a price of anarchy $1 - c/e$ in resource allocation games with nonnegative, nondecreasing concave welfare functions with maximum curvature $c \in [0, 1]$. In cases where the maximum curvature is not known, the guarantee corresponding to $c = 1$ still applies. Furthermore, we show that the local utility functions can be computed in polynomial time as nonnegative linear combination over a restricted set of functions with explicit expressions.

In the example we studied in Section 1.2, we observed that the price of anarchy achieved by the universal utility mechanism is near-optimal within sets of games induced by specialized welfare sets. Considering the gains in tractability and generality when using this mechanism, this small decrease in equilibrium efficiency guarantees may be acceptable. Future work should characterize the difference between the price of anarchy achieved by the universal utility mechanism and the best achievable price of anarchy within the set of games induced by a given set of welfare functions.

We observed that, in certain cases, the price of anarchy guarantees that we obtain match the best-achievable approximation ratios among polynomial-time centralized algorithms [4, 32]. An investigation into the potential connections between the best achievable price of anarchy in resource allocation games and the best achievable approximation ratio among polynomial-time centralized algorithms would reflect on the relative performance of distributed and centralized multiagent coordination algorithms.

Since the price of anarchy is a measure for the worst-case equilibrium efficiency within a family of instances, it may not be representative of the expected performance of a distributed algorithm designed using the game theoretic approach. This is demonstrated, for example, by the simulation results studied in Section 3. A relevant research direction is the design of player utility functions with the objective of maximizing the expected equilibrium efficiency.

ACKNOWLEDGMENTS

This work is supported by ONR Grant #N00014-20-1-2359 and AFOSR Grant #FA9550-20-1-0054. This research was, in part, funded by the U.S. Government. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

REFERENCES

- [1] Nodir Adilov, Peter J Alexander, and Brendan M Cunningham. 2015. An economic analysis of earth orbit pollution. *Environmental and Resource Economics* 60, 1 (2015), 81–98.
- [2] Gürdal Arslan, Jason R. Marden, and Jeff S. Shamma. 2007. Autonomous Vehicle-Target Assignment: A Game-Theoretical Formulation. *Journal of Dynamic Systems, Measurement, and Control* 129, 5 (2007), 584–596.
- [3] Siddharth Barman, Omar Fawzi, and Paul Fermé. 2020. Tight approximation bounds for concave coverage problems. *arXiv preprint arXiv:2010.00970* (2020).
- [4] Siddharth Barman, Omar Fawzi, Suprovat Ghoshal, and Emirhan Gürpınar. 2020. Tight approximation bounds for maximum multi-coverage. In *International Conference on Integer Programming and Combinatorial Optimization*. Springer, 66–77.
- [5] Vittorio Bilò and Cosimo Vinci. 2019. Dynamic taxes for polynomial congestion games. *ACM Transactions on Economics and Computation (TEAC)* 7, 3 (2019), 1–36.
- [6] Ioannis Caragiannis. 2013. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica* 66, 3 (2013), 512–540.
- [7] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. 2010. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms (TALG)* 7, 1 (2010), 1–31.
- [8] Rahul Chandan, Dario Paccagnan, and Jason R Marden. 2019. When smoothness is not enough: toward exact quantification and optimization of the price-of-anarchy. In *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 4041–4046.
- [9] George Christodoulou, Elias Koutsoupias, and Akash Nanavati. 2004. Coordination mechanisms. In *International Colloquium on Automata, Languages, and Programming*. Springer, 345–357.
- [10] Serdar Çolak, Antonio Lima, and Marta C González. 2016. Understanding congested travel in urban areas. *Nature communications* 7, 1 (2016), 1–8.
- [11] Michele Conforti and Gérard Cornuéjols. 1984. Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. *Discrete applied mathematics* 7, 3 (1984), 251–274.
- [12] Chong Ding, Bi Song, Akshay Morye, Jay A Farrell, and Amit K Roy-Chowdhury. 2012. Collaborative sensing in a distributed PTZ camera network. *IEEE Transactions on Image Processing* 21, 7 (2012), 3282–3295.
- [13] Uriel Feige. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)* 45, 4 (1998), 634–652.
- [14] Stephen Forbes. 2020. Blackjack.
- [15] Martin Gairing. 2009. Covering games: Approximation through non-cooperation. In *International Workshop on Internet and Network Economics*. Springer, 184–195.
- [16] Alberto Giubilini. 2019. Antibiotic resistance as a tragedy of the commons: An ethical argument for a tax on antibiotic use in humans. *Bioethics* 33, 7 (2019), 776–784.
- [17] Garrett Hardin. 1968. The Tragedy of the Commons. *Sci* 162, 3859 (1968), 1243–1248.
- [18] Sergiu Hart and Andreu Mas-Colell. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 5 (2000), 1127–1150.
- [19] Dorit S Hochbaum. 1996. Approximating covering and packing problems: set cover, vertex cover, independent set, and related problems. In *Approximation algorithms for NP-hard problems*. 94–143.
- [20] Nicole Immorlica, Li Erran Li, Vahab S Mirrokni, and Andreas S Schulz. 2009. Coordination mechanisms for selfish scheduling. *Theoretical computer science* 410, 17 (2009), 1589–1598.
- [21] Elias Koutsoupias and Christos Papadimitriou. 2009. Worst-case equilibria. *Computer science review* 3, 2 (2009), 65–69.
- [22] William Forster Lloyd. 1833. *Two lectures on the checks to population*.
- [23] Jason R Marden and Tim Roughgarden. 2014. Generalized efficiency bounds in distributed resource allocation. *IEEE Trans. Automat. Control* 59, 3 (2014), 571–584.
- [24] Jason R Marden and Adam Wierman. 2013. Distributed welfare games. *Operations Research* 61, 1 (2013), 155–168.
- [25] Robert A Murphey. 2000. Target-based weapon target assignment problems. In *Nonlinear assignment problems*. Springer, 39–53.
- [26] Dario Paccagnan, Rahul Chandan, Bryce L Ferguson, and Jason R Marden. 2020. Incentivizing efficient use of shared infrastructure: Optimal tolls in congestion games. *arXiv preprint arXiv:1911.09806v3* (2020).
- [27] Dario Paccagnan, Rahul Chandan, and Jason R Marden. 2019. Utility Design for Distributed Resource Allocation-Part I: Characterizing and Optimizing the Exact Price of Anarchy. *IEEE Trans. Automat. Control* (2019).
- [28] T. Roughgarden. 2014. Barriers to Near-Optimal Equilibria. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. 71–80.
- [29] Tim Roughgarden. 2015. Intrinsic robustness of the price of anarchy. *Journal of the ACM (JACM)* 62, 5 (2015), 1–42.
- [30] Walid Saad, Zhu Han, H Vincent Poor, and Tamer Basar. 2012. Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications. *IEEE Signal Processing Magazine* 29, 5 (2012), 86–105.
- [31] Jeff Shamma. 2008. *Cooperative control of distributed multi-agent systems*. John Wiley & Sons.
- [32] Maxim Sviridenko, Jan Vondrák, and Justin Ward. 2017. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research* 42, 4 (2017), 1197–1218.
- [33] Vasilis Syrgkanis and Eva Tardos. 2013. Composable and efficient mechanisms. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 211–220.
- [34] Adrian Vetta. 2002. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings*. IEEE, 416–425.

A PROOF OF THEOREM 1

The proof relies on a result in Chandan et al. [8], which we detail in the following proposition for the reader's convenience:

Proposition 1 (Thm. 6 [8]). *Consider the set of resource allocation games with a maximum of n players induced by local welfare functions W^1, \dots, W^m . Let (F^j, ρ^j) , $j = 1, \dots, m$, be solutions to m linear programs of the form⁵*

$$\begin{aligned} \min_{F, \rho} \quad & \rho \\ \text{s.t.} \quad & W^j(y) - \rho W^j(x) + (x-z)F(x) - (y-z)F(x+1) \leq 0, \\ & \forall (x, y, z) \in \mathcal{I}(n). \end{aligned} \quad (12)$$

Then, the local utility functions F^1, \dots, F^m maximize the price of anarchy and the corresponding price of anarchy is $\text{PoA}(\mathcal{G}) = \min_{j \in \{1, \dots, m\}} \frac{1}{\rho^j}$.

A.1 Proof of Lemma 1

PROOF. We first dispense with the situation where $n \leq \beta$. In this case, the local welfare function is identical to $W(x) = x$, and thus the price of anarchy is 1 for choice of $F(x) = W(x)/x$. For the remainder of the proof, we only consider $n > \beta$.

The remainder of the proof is structured as follows: (i) we introduce a relaxation of the linear program in Equation (12); (ii) in this relaxed linear program, we determine what are the most restrictive constraints for each $x \in \{1, \dots, n-1\}$; (iii) we show that a feasible solution to the relaxed linear program is nonincreasing, i.e., $F(x+1) \leq F(x)$, for every $x \in \{1, \dots, n-1\}$ such that $F(x) > 1 - \alpha$, and $F(x+1) = 1 - \alpha$ otherwise; (iv) we show that (F, ρ) as defined in the claim is a solution to the relaxed linear program for $n \rightarrow \infty$; and (v) we observe that (F, ρ) as defined in the claim is feasible in the linear program in Equation (12) and thus a solution to this linear program as well.

Relaxed linear program. First we consider a relaxation of the linear program in Equation (12). In this relaxed linear program, only the constraints where $z = \min\{0, x + y - n\}$ and $x, y \in \{0, \dots, n\}$ are retained. Finally, we exclude the constraint with $y = 0$, for all $x \in \{1, \dots, n-1\}$, resulting in the following *relaxed linear program*:

$$\begin{aligned} \max_{F \in \mathbb{R}^n, \rho \in \mathbb{R}} \quad & \rho \quad \text{subject to:} \\ W(y) - \rho W(x) + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1) & \leq 0, \\ \forall (x, y) \in \{0, \dots, n\} \times \{1, \dots, n\} \cup (n, 0). & \end{aligned} \quad (13)$$

Tightest constraints on ρ . We characterize what value $y \in \{1, \dots, n\}$ parameterizes the tightest constraint for each $x \in \{1, \dots, n-1\}$. For any $x = 1, \dots, n-1$, if $1 \geq F(x)$, $F(x+1) \geq 1 - \alpha$, we observe that the constraint with $y = \beta$ is strictest. For $y < \beta$, it holds that

$$\begin{aligned} \rho W(x) & \geq \beta + \min\{x, n-\beta\}F(x) - \min\{\beta, n-x\}F(x+1) \\ & \geq y + \min\{x, n-y\}F(x) - \min\{y, n-x\}F(x+1), \end{aligned}$$

where the final inequality holds when $x \leq n - \beta$ because $\beta - y \geq (\beta - y)F(x+1)$; when $n - \beta < x \leq n - y$ because $\beta - y - (x + \beta -$

⁵ $\mathcal{I}(n)$ is defined as the set of all triplets $(x, y, z) \in \{0, 1, \dots, n\}^3$ that satisfy: (i) $1 \leq x + y - z \leq n$ and $z \leq \min\{x, y\}$; and, (ii) $x + y - z = n$ or $(x - z)(y - z)z = 0$.

$n)F(x) \geq n - x - y \geq (n - x - y)F(x+1)$ since $x + \beta - n > 0$; and when $n - y < x$ because $\beta - y \geq (\beta - y)F(x)$. For constraints with $y > \beta$,

$$\begin{aligned} \rho W(x) & \geq \alpha\beta + (1 - \alpha)\beta + \min\{x, n - \beta\}F(x) - \min\{\beta, n - x\}F(x+1) \\ & \geq \alpha\beta + (1 - \alpha)y + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1), \end{aligned}$$

where the final inequality holds when $x \leq n - y$ because $(y - \beta)F(x+1) \geq (y - \beta)(1 - \alpha)$, when $n - y < x \leq n - \beta$ because $(1 - \alpha)(\beta - y) + (x + y - n)F(x) \geq (1 - \alpha)(\beta + x - n) \geq (\beta + x - n)F(x+1)$ since $x + y - n > 0 \geq \beta + x - n$, and when $n - \beta < x$ because $(1 - \alpha)(\beta - y) \geq (\beta - y)F(x)$.

For any $x = 1, \dots, n-1$, if $F(x) \geq 1 - \alpha \geq F(x+1)$ and $n - x \geq \beta$, then the constraint with $y = n - x$ is strictest among all constraints as, for any $y \neq n - x$, it holds that

$$\begin{aligned} \rho W(x) & \geq \alpha\beta + (1 - \alpha)(n - x) + xF(x) - (n - x)F(x+1) \\ & \geq \alpha\beta + (1 - \alpha)y + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1) \\ & \geq W(y) + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1), \end{aligned}$$

where the inequality holds because $(1 - \alpha)(n - x - y) \geq (n - x - y)F(x+1)$ when $x \leq n - y$ and $(1 - \alpha)(n - x - y) \geq (n - x - y)F(x)$ when $x > n - y$. For any $x = 1, \dots, n-1$, if $F(x) \geq 1 - \alpha \geq F(x+1)$ and $n - x < \beta$, then $y = \beta$ is strictest as for any $y \neq \beta$, it holds that

$$\begin{aligned} \rho W(x) & \geq \beta + \min\{x, n - \beta\}F(x) - \min\{\beta, n - x\}F(x+1) \\ & \geq W(y) + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1), \end{aligned}$$

where $\beta - y + (n - \beta - x)F(x) \geq (\beta - y)[1 - F(x)] + (n - x - y)F(x) \geq (n - x - y)F(x+1)$ when $x \leq n - y$ since $y \leq n - x < \beta$, $(\beta - y)[1 - F(x)] \geq 0$ when $x > n - y$ and $n - x < y \leq \beta$, and $(1 - \alpha)(\beta - y) + (y - \beta)F(x) \geq 0$ when $y > \beta$ since $F(x) \geq 1 - \alpha$.

For any $x = 1, \dots, n-1$, if $F(x+1), F(x) \leq 1 - \alpha$, then the constraint with $y = n$ is strictest among all constraints as, for any $y < n$, it holds that

$$\begin{aligned} \rho W(x) & \geq \alpha\beta + (1 - \alpha)n + \min\{x, 0\}F(x) - \min\{n, n - x\}F(x+1) \\ & \geq \alpha\beta + (1 - \alpha)y + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1) \\ & \geq W(y) + \min\{x, n - y\}F(x) - \min\{y, n - x\}F(x+1), \end{aligned}$$

where the second last inequality holds because $(n - y)[1 - \alpha - F(x+1)] \geq x[F(x) - F(x+1)]$ when $x \leq n - y$ and $(n - y)(1 - \alpha) \geq (n - y)F(x)$ when $x > n - y$.

Thus, if $F(x) \geq F(x+1) \geq 1 - \alpha$, it is sufficient to consider only the constraint with $y = \beta$ and $z = \max\{0, x + \beta - n\}$. If $F(x) \geq 1 - \alpha \geq F(x+1)$ and $n - x > \beta$, it is sufficient to consider only the constraint with $y = n - x$ and $z = 0$. Otherwise, if $F(x) \geq 1 - \alpha \geq F(x+1)$ and $n - x < \beta$, then we consider only the constraint $y = \beta$ and $z = \max\{0, x + y - n\}$. Finally, if $1 - \alpha \geq F(x) \geq F(x+1)$, then the constraint with $y = n$ and $z = x$ is the strictest.

Proof that a solution to Equation (13) has F 'nonincreasing'. For this portion of the proof, consider a function F defined for any given $\rho > 1$ as follows: $F(1) = W(1)$ and, for all $x \in \{1, \dots, n-1\}$,

$$F(x+1) = \max_{y \in \{1, \dots, n\}} \frac{\min\{x, n - y\}F(x) - W(x)\rho + W(y)}{\min\{y, n - x\}}.$$

For conciseness, we will use the shorthand $\kappa_x = \frac{\min\{x, n-y^*\}}{\min\{y^*, n-x\}}$, $\lambda_x = \frac{W(x)}{\min\{y^*, n-x\}}$ and $\mu_x = \frac{W(y^*)}{\min\{y^*, n-x\}}$ where $y^* \in \{1, \dots, n\}$ maximizes the above expression for each $x \in \{1, \dots, n\}$. Thus, $F(x+1) = \kappa_x F(x) - \lambda_x \rho + \mu_x$ for each $x \in \{1, \dots, n\}$.

We assume, by contradiction, that we are given ρ such that F is increasing at some index, i.e., $F(x+1) > F(x)$ for $x \in \{1, \dots, n\}$. Lemma 3 in Appendix A.2 shows that, if this is the case, then F must continue to increase, so that $F(n) > F(n-1)$. We wish to show the following: (i) if F first increases at a point $x \in \{1, \dots, n-1\}$ where $F(x) > 1 - \alpha$, this leads to a contradiction for the value of ρ ; and, (ii) if F first increases at a point x where $F(x) \leq 1 - \alpha$, then either $F(j) \leq 1 - \alpha$ for all $j \geq x$ is feasible or (F, ρ) is infeasible. It is important to note that the value $nF(n)/W(n)$ must be bounded, otherwise

$$\rho \geq \max_{y \in \{0, \dots, n\}} \frac{W(y) + (n-y)F(n)}{W(n)} \geq \frac{nF(n)}{W(n)} \rightarrow \infty.$$

This is a contradiction as the price of anarchy will be at least 0.5, even if we use the marginal contribution utility [34]. Since we are optimizing for the price of anarchy (i.e., $1/\rho$), we need only consider values of ρ no greater than 2.

Observe that if F first increases at some point $x \in \{1, \dots, n\}$ such that $F(x) > 1 - \alpha$, then $F(n) > F(n-1) > 1 - \alpha$ and

$$\begin{aligned} \rho &= \max_{y \in \{0, \dots, n\}} \frac{W(y) + (n-y)F(n)}{W(n)} \\ &\geq \max_{y \in \{1, \dots, n\}} \frac{W(y) + (n-y)F(n)}{W(n)} \\ &= \max_{y \in \{1, \dots, n-1\}} \frac{W(y) + (n-y)F(n)}{W(n)} \\ &> \max_{y \in \{1, \dots, n-1\}} \frac{W(y) + (n-y)F(n-1)}{W(n)}, \end{aligned}$$

where the first inequality holds because we reduce the domain of maximization, the second equality holds because $W(n-1) + F(n) > W(n)$ since $n > \beta$ and $F(n) > 1 - \alpha$, and the final inequality holds because $F(n) > F(n-1)$. Since $y \leq \beta \leq n-1$ corresponds with the strictest constraints when $F(n), F(n-1) \geq 1 - \alpha$, we can substitute

$$\max_{y \in \{1, \dots, n-1\}} [W(y) + (n-y)F(n-1)] = F(n) + \rho W(n-1)$$

in the former bound on ρ to get $\rho > [F(n) + \rho W(n-1)]/W(n)$. Since $n > \beta$, this implies that

$$\begin{cases} F(n) < 0 & \text{if } \alpha = 1, \\ \rho > \frac{F(n)}{W(n)-W(n-1)} = \frac{F(n)}{1-\alpha} & \text{if } \alpha \in [0, 1) \end{cases}$$

For $\alpha = 1$ this is a contradiction, since we have that $F(n) > 1 - \alpha = 0$. For the remaining $\alpha \in [0, 1)$, we want to prove that this also gives rise to a contradiction. To do so, we show that

$$\frac{F(n)}{1-\alpha} \geq \max_{y \in \{0, \dots, n\}} \frac{W(y) + (n-y)F(n)}{W(n)} = \rho.$$

Observe that $y = \beta$ maximizes the right-hand side if $1 - \alpha < F(n) \leq 1$, and $y = 0$ maximizes the right-hand side if $F(n) > 1$. For

$F(n) \in (1 - \alpha, 1]$ and $y = \beta$, it holds that

$$\begin{aligned} \frac{F(n)}{1-\alpha} - \frac{W(\beta) + (n-\beta)F(n)}{W(n)} &\geq 0 \\ \iff F(n)W(n) - (1-\alpha)[W(\beta) + (n-\beta)F(n)] &\geq 0 \\ \iff F(n) \frac{W(n) - W(\beta)}{n-\beta} &= (1-\alpha)F(n) \end{aligned}$$

where the first and second line are equivalent because $\alpha \in [0, 1)$ and the third line implies the second because $F(n) > 1 - \alpha$. The final inequality holds because $[W(n) - W(\beta)] = (n - \beta)(1 - \alpha)$ and $n > \beta$. For $F(n) > 1$ and $y = 0$, it holds that

$$\frac{F(n)}{1-\alpha} - \frac{nF(n)}{W(n)} \geq 0$$

since $W(n)/n \geq 1 - \alpha$ by definition. Thus, in the above reasoning, we have shown that, if F first increases at a point $x \in \{1, \dots, n\}$ and $F(x) > 1 - \alpha$, then, if $\alpha = 1$, it holds that $1 - \alpha < F(x) < \dots < F(n) < 1 - \alpha$; and, if $\alpha \in [0, 1)$, it holds that

$$\rho > \frac{F(n)}{1-\alpha} \geq \max_{y \in \{0, \dots, n\}} \frac{W(y) + (n-y)F(n)}{W(n)} = \rho,$$

which is a contradiction.

Now we consider the scenario where we are at a point x such that $F(x) \leq 1 - \alpha$ and F is monotonically nonincreasing before x . We show that either selecting $F(x) = \dots = F(n) = 1 - \alpha$ is feasible for ρ or that the value ρ is infeasible. We first consider the case where the strictest constraint on the value of $F(x+1)$ has $y \leq n-x$ and show that $F(x+1)$ cannot be greater than $F(x) \leq 1 - \alpha$. In the proof of Lemma 3 in Appendix A.2, we showed that if $y \leq n-x$, $F(x+1) > F(x)$ and $F(x) \leq F(x-1)$, then it must hold that $F(x) > [W(x) - W(x-1)]\rho \geq 1 - \alpha$. As we have assumed $F(x) \leq 1 - \alpha$, it must be that $F(x+1) \leq F(x) \leq 1 - \alpha$ if $y \leq n-x$. We complete our reasoning for the case when $y > n-x$ corresponds to the strictest constraint on the value of $F(x+1)$. We showed above that if $F(x) \leq 1 - \alpha$ and $x > n-y$, then the strictest constraint is parameterized by $y = n$. For any $x \geq \beta$, it must hold that $F(x+1) \geq -W(x)\rho/(n-x) + W(n)/(n-x)$. Since $-W(x)\rho/(n-x) + W(n)/(n-x) \leq 1 - \alpha$ the constraint is satisfied for choice of $F(x+1) \leq 1 - \alpha$. Else, if $x < \beta$, since $\beta < n$, $F(x+1) > 1 - \alpha$ implies that $F(n) > F(n-1) > 1 - \alpha$, since $n-1 \geq \beta$. We already proved above that this scenario leads to a contradiction on the value of ρ . Repeating this reasoning for all $j > x$ such that $F(x) \leq 1 - \alpha$, we argue that $F(j) = 1 - \alpha$ is feasible. Since the strictest constraint for each $F(j), F(j+1) \leq 1 - \alpha$ has $y = n$, there is no recursion and the optimal value ρ has no dependence on the values of $F(x), \dots, F(n)$, even if it begins increasing. We have also shown that the lower bound on F is lower than or equal to $1 - \alpha$ for any feasible ρ , and so, F with $F(j) = 1 - \alpha$, for all $j \in \{x, \dots, n\}$, must be feasible.

For any feasible ρ , we have successfully shown that $F(x)$ must be nonincreasing when it is greater than $1 - \alpha$, and that $F(x) = \dots = F(n) = 1 - \alpha$ is feasible otherwise. This concludes this part of the proof.

Proof that (F, ρ) solves Equation (13). We begin by showing that (F, ρ) as defined in the claim are feasible. For $x = 0$, the constraints

in Equation (13) read as $F(1) \geq W(y)/\min\{y, n-x\}$, for all $y = 1, \dots, n$, which is satisfied for $F(1) = W(1)$. Now consider $(x, y) \in \{1, \dots, n-1\} \times \{1, \dots, n\}$. In the above reasoning, we showed that a feasible (F, ρ) within Equation (13) will have F nonincreasing while $F(x) > 1 - \alpha$ and $F(x) = 1 - \alpha$ otherwise. Furthermore, we showed that when $F(x) \geq F(x+1) > 1 - \alpha$ or when $F(x) \geq 1 - \alpha \geq F(x+1)$ and $n-x < \beta$, then the strictest constraint has $y = \beta$. Observe that $\kappa_x = \min\{x, n-\beta\}/\min\{\beta, n-x\}$, $\lambda_x = W(x)/\min\{\beta, n-x\}$ and $\mu_x = \beta/\min\{\beta, n-x\}$ correspond with the recursive definition of $F(x+1)$.

We showed above that $F(x+1) = 1 - \alpha$ when $\kappa_x F(x) - \lambda_x \rho + \mu_x \leq 1 - \alpha$, is feasible as long as ρ is feasible, since the values of F less than or equal to $1 - \alpha$ have no impact on the optimal value of the relaxed linear program. Consider the expression for ρ that can be obtained by completing the recursion as follows,

$$1 - \alpha = F(\hat{x}+1) \geq \Pi_{u=1}^{\hat{x}} \kappa_u F(1) + \sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v) (\mu_u - \lambda_u \rho) + \mu_{\hat{x}} - \lambda_{\hat{x}} \rho.$$

Rearranging this expression, we obtain,

$$\rho \geq \frac{\Pi_{u=1}^{\hat{x}} \kappa_u F(1) + \sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v) \mu_u + \mu_{\hat{x}} + \alpha - 1}{\sum_{u=1}^{\hat{x}-1} (\Pi_{v=u+1}^{\hat{x}} \kappa_v) \lambda_u + \lambda_{\hat{x}}}.$$

Observe that for $n \rightarrow \infty$, $\min\{x, n-\beta\} = x$ and $\min\{\beta, n-x\} = \beta$. Thus, the above expression for ρ simplifies to

$$\begin{aligned} \rho &\geq \frac{\frac{\hat{x}!}{\beta^{\hat{x}}} + \sum_{u=1}^{\hat{x}-1} \frac{\hat{x}!}{j! \beta^{\hat{x}-j}} \frac{1}{\beta^{\hat{x}-j}} + 1 + \alpha - 1}{\sum_{j=1}^{\hat{x}} \frac{\hat{x}!}{j! \beta^{\hat{x}-j}} \frac{j}{\beta} + \sum_{j=\beta+1}^{\hat{x}-1} \frac{\hat{x}!}{j! \beta^{\hat{x}-j}} \frac{\alpha \beta + (1-\alpha)j}{\beta}} \\ &= \frac{1 + \sum_{j=1}^{\hat{x}-1} \frac{\beta^j}{j!} + \alpha \frac{\beta^{\hat{x}}}{\hat{x}!}}{\sum_{j=0}^{\hat{x}-1} \frac{\beta^j}{j!} + \sum_{j=\beta+1}^{\hat{x}-1} \frac{\beta^j}{j!} \frac{\alpha \beta + (1-\alpha)j}{\beta}} = \frac{e^\beta}{e^\beta - \alpha \frac{\beta^\beta}{\beta!}}. \end{aligned}$$

Noting that $\text{PoA} = 1/\rho$ concludes this part of the proof.

Feasibility of (F, ρ) in Equation (12). To conclude the proof, we simply observe that since $F(x)$ is nonincreasing for all x , the strictest constraints in the linear program in Equation (12) correspond with the choice of $z = \min\{0, x + y - n\}$. Thus, since (F, ρ) is a solution to the relaxed linear program and feasible in the original linear program, it must also be a solution to the original. \square

A.2 Proof of Lemma 3

Lemma 3. *Let W be a nonnegative, nondecreasing concave function, and let $\rho \geq 1$ be a given parameter. Further, define the function F such that $F(1) = W(1)$ and*

$$F(j+1) := \max_{\ell \in \{1, \dots, n\}} \frac{\min\{j, n-\ell\} F(j) - W(j) \rho + W(\ell)}{\min\{\ell, n-j\}}, \quad (14)$$

for all $j = 1, \dots, n-1$. Then, for the lowest value $\hat{j} = 1, \dots, n-1$ such that $F(\hat{j}+1) > F(\hat{j})$, it must hold that $F(j+1) > F(j)$ for all $j = \hat{j}, \dots, n-1$.

PROOF. The proof is presented in two parts as follows: in part (i), we identify an inequality that must hold given that $F(\hat{j}+1) > F(\hat{j})$ for $1 \leq \hat{j} \leq n-1$ as defined in the claim; and, in part (ii), we use

a recursive argument to prove that $F(j+1) > F(j)$ holds for all $\hat{j}+1 \leq j \leq n-1$, using the inequality we derived in part (i).

Part (i). We define ℓ_j^* as one of the arguments that minimizes the right-hand side of Equation (14) for each $x = 1, \dots, n-1$. By assumption, it must hold that $F(\hat{j}+1) > F(\hat{j})$, which implies that

$$\begin{aligned} F(\hat{j}) &< \max_{1 \leq \ell \leq n} \frac{\min\{\hat{j}, n-\ell\}}{\min\{\ell, n-\hat{j}\}} F(\hat{j}) - \frac{W(\hat{j})}{\min\{\ell, n-\hat{j}\}} \rho + \frac{W(\ell)}{\min\{\ell, n-\hat{j}\}} \\ &= \max_{1 \leq \ell \leq n} \frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}} \rho \\ &\quad + \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}} + \frac{\min\{\hat{j}, n-\ell\}}{\min\{\ell, n-\hat{j}\}} F(\hat{j}) \\ &\quad - \frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j})}{\min\{\ell, n-\hat{j}\}} \rho \\ &\quad + \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}} \rho + \frac{W(\ell)}{\min\{\ell, n-\hat{j}\}} \\ &\quad - \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}}, \end{aligned}$$

where the strict inequality holds by definition of $F(\hat{j}+1)$. Recall that

$$F(\hat{j}) := \max_{1 \leq \ell \leq n} \frac{\min\{\hat{j}-1, n-\ell\}}{\min\{\ell, n-\hat{j}+1\}} F(\hat{j}-1) - \frac{W(\hat{j}-1)}{\min\{\ell, n-\hat{j}+1\}} \rho + \frac{W(\ell)}{\min\{\ell, n-\hat{j}+1\}}.$$

Thus, if $\ell_j^* \leq n - \hat{j}$, the above strict inequality with $F(\hat{j})$ can only be satisfied if

$$F(\hat{j}+1) > F(\hat{j}) \geq \hat{j} F(\hat{j}) - (\hat{j}-1) F(\hat{j}-1) > [W(\hat{j}) - W(\hat{j}-1)] \cdot \rho.$$

Similarly, if $\ell_j^* \geq n - \hat{j} + 1$, then it must hold that

$$\begin{aligned} (n - \ell_j^*) &\left[\frac{F(\hat{j})}{n - \hat{j}} - \frac{F(\hat{j}-1)}{n - \hat{j} + 1} \right] + \left[\frac{1}{n - \hat{j}} - \frac{1}{n - \hat{j} + 1} \right] W(\ell_j^*) \\ &> \left[\frac{W(\hat{j})}{n - \hat{j}} - \frac{W(\hat{j}-1)}{n - \hat{j} + 1} \right] \cdot \rho \\ &\implies \left[\frac{1}{n - \hat{j}} - \frac{1}{n - \hat{j} + 1} \right] [(n - \ell_j^*) F(\hat{j}) + W(\ell_j^*)] \\ &> \left[\frac{W(\hat{j})}{n - \hat{j}} - \frac{W(\hat{j}-1)}{n - \hat{j} + 1} \right] \cdot \rho \\ &\iff F(\hat{j}+1) > [W(\hat{j}) - W(\hat{j}-1)] \rho, \end{aligned}$$

where the first line implies the second line because $F(\hat{j}) \leq F(\hat{j}-1)$, by the definition of \hat{j} in the claim, and the second line is equivalent to the third by the definitions of $F(\hat{j}+1)$ and ℓ_j^* . This concludes part (i) of the proof.

Part (ii). In this part of the proof, we show by recursion that if $F(\hat{j}+1) > F(\hat{j})$, then $F(j+1) > F(j)$ for all $j = \hat{j}+1, \dots, n-1$. We do so by showing that, if $F(j) > F(j-1) > \dots > F(\hat{j}+1)$ for any $\hat{j}+1 \leq j \leq n-1$, then it must hold that $F(j+1) > F(j)$. Thus, in the following reasoning, we assume that $\hat{j}+1 \leq j \leq n-1$, and that $F(j) > F(j-1) > \dots > F(\hat{j}+1)$.

We begin with the scenario in which $\ell_{j-1}^* < n - j + 1$, which gives us that $\ell_{j-1}^* \leq n - j$. Recall that

$$F(j+1) := \max_{1 \leq \ell_j \leq n} \frac{\min\{j, n - \ell_j\}}{\min\{\ell_j, n - j\}} F(j+1) - \frac{W(j)}{\min\{\ell_j, n - j\}} \rho + \frac{W(\ell_j)}{\min\{\ell_j, n - j\}}.$$

Thus, it must hold that

$$\begin{aligned} & F(j+1) \\ &= \max_{1 \leq \ell_j \leq n} \frac{\min\{j-1, n - \ell_j\}}{\min\{\ell_j, n - j + 1\}} F(j-1) - \frac{W(j-1)}{\min\{\ell_j, n - j + 1\}} \rho \\ & \quad + \frac{W(\ell_j)}{\min\{\ell_j, n - j + 1\}} + \frac{\min\{j, n - \ell_j\}}{\min\{\ell_j, n - j\}} F(j+1) \\ & \quad - \frac{\min\{j-1, n - \ell_j\}}{\min\{\ell_j, n - j + 1\}} F(j-1) - \frac{W(j)}{\min\{\ell_j, n - j\}} \rho \\ & \quad + \frac{W(j-1)}{\min\{\ell_j, n - j + 1\}} \rho + \frac{W(\ell_j)}{\min\{\ell_j, n - j\}} \\ & \quad - \frac{W(\ell_j)}{\min\{\ell_j, n - j + 1\}} \\ & \geq F(j) + \frac{j}{\ell_{j-1}^*} F(j) - \frac{j-1}{\ell_{j-1}^*} F(j-1) - \frac{W(j) - W(j-1)}{\ell_{j-1}^*} \rho \\ & > F(j) + \frac{1}{\ell_{j-1}^*} F(\hat{j}+1) - \frac{1}{\ell_{j-1}^*} [W(j) - W(j-1)] \rho \\ & > F(j), \end{aligned}$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell_{j-1}^*$, the second inequality holds because $F(j) > F(j-1)$ and $F(j) \geq F(\hat{j}+1)$, by assumption, and the final inequality holds by the identity we showed in part (i) and because $W(\cdot)$ is concave.

Next, consider the scenario in which $\ell_{j-1}^* > n - j + 1$. Observe that

$$\begin{aligned} & F(j+1) \\ & \geq F(j) + (n - \ell_{j-1}^*) \left[\frac{F(j)}{n-j} - \frac{F(j-1)}{n-j+1} \right] \\ & \quad + \left[\frac{1}{n-j} - \frac{1}{n-j+1} \right] W(\ell_{j-1}^*) - \frac{W(j)}{n-j} \rho + \frac{W(j-1)}{n-j+1} \rho \\ & > F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1} \right] [(n - \ell_{j-1}^*) F(j-1) + W(\ell_{j-1}^*)] \\ & \quad - \frac{W(j)}{n-j} \rho + \frac{W(j-1)}{n-j+1} \rho \\ & = F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1} \right] [(n-j+1) F(j) + W(j-1) \rho] \\ & \quad - \frac{W(j)}{n-j} \rho + \frac{W(j-1)}{n-j+1} \rho \\ & \geq F(j) + \frac{1}{n-j} F(\hat{j}+1) - \frac{1}{n-j} [W(j) - W(j-1)] \rho \\ & > F(j), \end{aligned}$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell_{j-1}^*$, the second inequality holds because $F(j) > F(j-1)$, by assumption, the equality holds by the definitions of $F(j)$ and ℓ_{j-1}^* ,

the third inequality holds because $F(j) \geq F(\hat{j}+1)$, by assumption, and the final inequality holds by the identity we showed in part (i) and because $W(\cdot)$ is concave.

Finally, we consider the scenario in which $\ell_{j-1}^* = n - j + 1$. Observe that

$$\begin{aligned} & F(j+1) \\ & \geq F(j) + \frac{j-1}{n-j} F(j) - \frac{j-1}{n-j+1} F(j-1) - \frac{W(j)}{n-j} \rho + \frac{W(j-1)}{n-j+1} \rho \\ & \quad + \left[\frac{1}{n-j} - \frac{1}{n-j+1} \right] W(n-j+1) \\ & > F(j) + \left[\frac{1}{n-j} - \frac{1}{n-j+1} \right] [(n - \ell_{j-1}^*) F(j-1) + W(\ell_{j-1}^*)] \\ & \quad - \frac{W(j)}{n-j} \rho + \frac{W(j-1)}{n-j+1} \rho \\ & > F(j), \end{aligned}$$

where the first inequality holds by evaluating the maximization at $\ell_j = \ell_{j-1}^*$, the second inequality holds because $F(j) > F(j-1)$, by assumption, and the final inequality holds by the same reasoning as for $\ell_{j-1}^* > n - j + 1$. \square

A.3 Proof of Lemma 2

PROOF. The proof is by construction. Define coefficients $\eta_1 := [2W(1) - W(2)]/c$, $\eta_j := [2W(j) - W(j-1) - W(j+1)]/c$, $j = 2, \dots, n-1$, and $\eta_n := W(1) - \sum_{j=1}^{n-1} \eta_j = W(1) - [W(1) + W(n-1) - W(n)]/c$. It is straightforward to verify that $\eta_j \geq 0$ for all $k = 1, \dots, n$ recalling that $W(0) = 0$ and $W(x)$ is nonnegative, nondecreasing concave for $x \geq 0$. We defer the proof that $W(x) = \sum_{k=1}^n \eta_k \cdot V_k^c(x)$ for all $x = 1, \dots, n$ to the proof of Corollary 1, where one need only substitute $W^{ub}(x) = x$ and $W^{lb}(x) = V_1^c(x)$, for $x \geq 0$. \square

B PROOF OF COROLLARY 1.

PROOF. First, observe that there must exist functions W^{ub} and W^{lb} . Simply observe that $W^{ub}(x) = x$ and $W^{lb}(x) = V_1^1(x) = \min\{x, 1\}$ are valid for any set of nonnegative, nondecreasing concave functions.

The rest of the proof follows by construction. Define the coefficients η_j , $j = 0, \dots, n$, as follows:

$$\eta_1 = \frac{W^{ub}(2) - W^{ub}(1) - W(2) + W(1)}{W^{ub}(2) - W^{ub}(1) - W^{lb}(2) + W^{lb}(1)},$$

$$\eta_j = \frac{W^{ub}(j+1) - W^{ub}(j) - W(j+1) + W(j)}{W^{ub}(j+1) - W^{ub}(j) - W^{lb}(j+1) + W^{lb}(j)} - \sum_{k=1}^{j-1} \eta_k,$$

for $j = 2, \dots, n-1$ and $\eta_n = 1 - \sum_{k=1}^{n-1} \eta_k$.

First, we prove that the coefficients η_1, \dots, η_n are nonnegative. It is simple to see that $\eta_1 \geq 0$ since $W^{ub}(2) - W^{ub}(1) \geq W(2) - W(1) \geq W^{lb}(2) - W^{lb}(1)$. Similarly, $\eta_n \geq 0$ since $\eta_n = 1 - [W^{ub}(n) - W^{ub}(n-1) - W(n) + W(n-1)]/[W^{ub}(n) - W^{ub}(n-1) - W^{lb}(n) + W^{lb}(n-1)]$.

1) - $W^{lb}(n) + W^{lb}(n-1)$]. Finally, for any $j \in \{2, \dots, n-1\}$,

$$\begin{aligned} \eta_j &= \frac{W^{ub}(j+1) - W^{ub}(j) - W(j+1) + W(j)}{W^{ub}(j+1) - W^{ub}(j) - W^{lb}(j+1) + W^{lb}(j)} \\ &\quad - \frac{W^{ub}(j) - W^{ub}(j-1) - W(j) + W(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \\ &\geq \frac{W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \\ &\quad - \frac{W(j+1) - 2W(j) + W(j-1)}{W^{ub}(j) - W^{ub}(j-1) - W^{lb}(j) + W^{lb}(j-1)} \geq 0, \end{aligned}$$

where the equality holds by definition, the first inequality holds because $W^{lb}(j+1) - 2W^{lb}(j) + W^{lb}(j-1) \geq W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1)$ and the final inequality holds because $W^{ub}(j+1) - 2W^{ub}(j) + W^{ub}(j-1) \geq W(j+1) - 2W(j) + W(j-1)$.

We conclude the proof by observing that, for all $x = 1, \dots, n$,

$$\begin{aligned} &\sum_{k=1}^{x-1} \eta_k W^{lb}(x) + \sum_{k=1}^{x-1} \eta_k [W^{ub}(k) - W^{lb}(k)] + \sum_{k=x}^n \eta_k W^{ub}(x) \\ &= \frac{W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1)}{W^{ub}(x) - W^{ub}(x-1) - W^{lb}(x) + W^{lb}(x-1)} W^{lb}(x) \\ &\quad + \sum_{k=1}^{x-1} \eta_k [W^{ub}(k) - W^{lb}(k)] \\ &\quad + \left[1 - \frac{W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1)}{W^{ub}(x) - W^{ub}(x-1) - W^{lb}(x) + W^{lb}(x-1)} \right] W^{ub}(x) \\ &= W^{ub}(x) - W^{ub}(x) - W^{ub}(x-1) - W(x) + W(x-1) \\ &\quad + \sum_{k=1}^{x-2} \eta_k [W^{ub}(k) - W^{lb}(k) - W^{ub}(x-1) + W^{lb}(x-1)] \\ &= W(x) + \sum_{k=1}^{x-2} \eta_k [W^{ub}(k) - W^{lb}(k) - W^{ub}(x-1) + W^{lb}(x-1)] \\ &\quad + W^{ub}(x-1) - W(x-1) \\ &= W(x), \end{aligned}$$

where the final equality holds once the expression is simplified for the remaining η_k values. \square