

Testing Core Membership in Public Goods Economies*

Greg Bodwin

MIT EECS, Cambridge, MA, USA
greg.bodwin@gmail.com

Abstract

This paper develops a recent line of economic theory seeking to understand public goods economies using methods of topological analysis. Our first main result is a very clean characterization of the economy’s core (the standard solution concept in public goods). Specifically, we prove that a point is in the core iff it is Pareto efficient, individually rational, and the set of points it dominates is path connected.

While this structural theorem has a few interesting implications in economic theory, the main focus of the second part of this paper is on a particular algorithmic application that demonstrates its utility. Since the 1960s, economists have looked for an efficient computational process that decides whether or not a given point is in the core. All known algorithms so far run in exponential time (except in some artificially restricted settings). By heavily exploiting our new structure, we propose a new algorithm for testing core membership whose computational bottleneck is the solution of $O(n)$ convex optimization problems on the utility function governing the economy. It is fairly natural to assume that convex optimization should be feasible, as it is needed even for very basic economic computational tasks such as testing Pareto efficiency. Nevertheless, even without this assumption, our work implies for the first time that core membership can be efficiently tested on (e.g.) utility functions that admit “nice” analytic expressions, or that appropriately defined ε -approximate versions of the problem are tractable (by using modern black-box ε -approximate convex optimization algorithms).

1998 ACM Subject Classification J.4 Social and Behavioral Sciences

Keywords and phrases Algorithmic Game Theory, Economics, Algorithms, Public Goods, Coalitional Stability

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.45

1 Introduction

1.1 Background on Public Goods Economics

A basic question in economics is to understand the forces governing the production of public goods. A good is *public* if its use by one person does not reduce its availability to others, and if none are excluded from using the good. Examples include public parks, research information, a clean environment, national defense, radio broadcasts, and so on.

Public goods economies were first explicitly abstracted in a classic paper by Samuelson in 1954 [24], and have since become central objects of study for economists. An important feature of public goods economies is that they are not well modeled by the individualistic “best-response dynamics” which govern familiar economic equilibrium concepts such as the Nash or Walrasian equilibrium. Rather, public goods typically arise as the result of a

* A full version of the paper is available at <https://arxiv.org/abs/1705.01570>.



© Greg Bodwin;

licensed under Creative Commons License CC-BY

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).

Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;

Article No. 45; pp. 45:1–45:14



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



communal process – e.g. negotiations, treaties, or taxes – that allow the cost of production to be amortized over all agents that stand to benefit from the good. Accordingly, public goods economics inspired the development of *cooperative game theory*, which seeks to understand these cooperative dynamics and when an agreement to produce public goods is “stable” or when it is doomed to fall apart.

What, exactly, does “stability” mean in this context? The most standard notion is *coalitional stability*, which is given as follows:

► **Definition (Informal).** Let \mathbf{a} be an outcome in a public goods economy with agents N (i.e. \mathbf{a} describes the amount of work each agent contributes to produce a public good). We say that \mathbf{a}' is a *deviation* on \mathbf{a} for a nonempty coalition of agents $C \subseteq N$ if no agents in $N \setminus C$ perform work (i.e. $a'_i = 0$ for all $i \notin C$) and all agents in C prefer \mathbf{a}' to \mathbf{a} . If there is no deviation on \mathbf{a} for any coalition, then we say \mathbf{a} is *coalitionally stable*. The set of coalitionally stable outcomes is called the *core* of the economy.

Since its inception in the late 1800’s [10],¹ the core and cooperative game theory in general have played major roles in many successful economic research programs. More background can be found in most modern game theory textbooks, e.g. [23, 8].

1.2 (Non-)Algorithmic Properties of Public Goods Economies

An inherent conceptual drawback of coalitional stability is its exponential-size definition. In other words, for an outcome \mathbf{a} to be coalitionally stable, *every single one* of the $2^n - 1$ possible coalitions of agents must not have a deviation. Hence, the naive algorithms for testing the coalitional stability of \mathbf{a} must perform computations for all $2^n - 1$ coalitions, and so they suffer exponential runtime. Coalitional stability is not a very convincing solution concept if its implicit notion of “instability” assumes that agents can quickly make exponential time computations in order to find deviations whenever they exist.

This problem has led economists studying public goods to seek more clever methods for solving computational problems related to the core of a public goods economy, which avoid this exponential behavior. Some of the initial work in this vein attacked the closely-related problem of simply outputting any core outcome. The first such solution appeared in the 1960s, when Scarf [25] proved that “balanced” games have nonempty cores, by means of an (exponential time) algorithm that outputs a core point and provably always terminates. Similar results were proved in the setting of public goods economies by Chander and Tulkens [5] and Elliot and Golub [12]. Meanwhile, followup work has suggested that the slow runtime of Scarf’s algorithm may be inherent to the problem: Kintali et al. [20] showed that Scarf’s algorithm cannot be improved to polynomial runtime (unless $P = PPAD$), and Deng and Papadimitriou [9] showed that it is NP Complete just to detect whether or not the core is empty (let alone find a point in the core), even in the simple class of graphical games, and even when Scarf’s assumption of “balance” is dropped. Other notable hardness results in this vein have come from Conitzer and Sandholm [7] and Greco et al. [18].

There has also been considerable prior work on the “membership testing” problem of determining whether a point taken on input is in the core (this is the question addressed in this paper). Deng and Papadimitriou’s work [9] also implies that membership testing is NP hard even in the restricted setting of graphical games, although there are straightforward

¹ Some of the early work on cooperative game theory used the name *contract curve* instead of core. The term *core* was coined in [15].

efficient algorithms in the further restricted setting where the game is superadditive. Conitzer and Sandholm [6] showed that membership testing is co-NP complete in games where coalition values have a “multiple-issue” representation in polynomial space. Faigle et al. [13] showed that the problem is NP complete in a variant of graphical games where payoffs are given by minimum spanning trees of subgraphs. Sung and Dimitrov [26] showed co-NP completeness for membership testing in “hedonic coalition formation games.” Goemans and Skutella showed NP completeness for both emptiness and membership testing in “facility location games,” and gave formulations of these problems as LP relaxations [16]. There has been work on games defined by *marginal contribution nets* (MC-nets) [19, 11], in which values attainable by coalitions are determined by succinct logical formulae. Li and Conitzer [22] studied emptiness testing and membership testing under various classes of formulae, and obtained various algorithms or NP hardness results depending on the complexity of the formulae allowed.

1.3 Our Results

A recently popular trend in public goods research has been to model economies as networks, and then seek to analyze the economy by studying the topological properties of the underlying “benefits network,” describing the ability of agents to transfer utility to each other at any given point (see for example [3, 4, 1, 2, 12]). Much of the initial work focused on Nash equilibria [3, 4, 1, 2] of the economy. A major stride was recently taken by Elliot and Golub [12], who extended the theory to show that the Lindahl equilibria² of an economy are precisely the points that are eigenvalues of their own benefits network. They further discuss connections between the Lindahl equilibria and the core – in particular, in any standard model (including theirs) all Lindahl equilibria are also core outcomes [14]. However, they raise an interesting open question to more precisely characterize the core [17] in a similar vein.

Our first main result achieves this goal. We show that the core can be characterized as follows:

► **Theorem 1.** *Let \mathbf{a} be an outcome in a public goods economy and let $\mathbf{D}_{\mathbf{a}}$ be the set of points that no agent prefers to \mathbf{a} . Then \mathbf{a} is in the core if and only if it is Pareto efficient, individually rational, and $\mathbf{D}_{\mathbf{a}}$ is path connected.*

(Here, the definition of path connectedness is the standard topological one: for any two points $\mathbf{x}, \mathbf{y} \in \mathbf{D}_{\mathbf{a}}$, there is a continuous function $f : [0, 1] \rightarrow \mathbb{R}_{\geq 0}^n$ satisfying $f(0) = \mathbf{x}$, $f(1) = \mathbf{y}$, and $f(\lambda) \in \mathbf{D}_{\mathbf{a}}$ for all $0 \leq \lambda \leq 1$. The image of $[0, 1]$ under f is called a path.)

An interesting consequence of this theorem is a precise description of the relationship between Lindahl equilibria and core outcomes:

► **Theorem 2.** *Assuming that the utility function \mathbf{u} is differentiable, the Lindahl equilibria of a public goods economy are precisely the core points whose core membership can be certified using only local information.*

The proof of this corollary is essentially immediate by combining Theorem 1 with a more technical phrasing of Elliot and Golub’s result.

While we believe that these two structural theorems hold intrinsic interest, the second half of this paper is intended to demonstrate their power by an application to the algorithmic

² The Lindahl equilibria of an economy are the competitive equilibria that would be reached if market externalities were truthfully reported and then bought and sold on an open market. A formal definition is not necessary to read this paper, but can be found in e.g. [23, 8].

problem discussed earlier. We have previously suggested the intuition that the algorithmic core membership testing problem is hard because the naive algorithms must check an exponential number of coalitions for a potential deviation. However, Theorem 1 lets us avoid this brute-force behavior: after checking for Pareto efficiency and individual rationality (which is quite easy), we are left only with the task of checking whether or not $\mathbf{D}_{\mathbf{a}}$ is path-connected. The complexity of this task is non-obvious, but we show that it can be done fairly efficiently, yielding the following result:

► **Theorem 3.** *Given an outcome \mathbf{a} in an n -agent public goods economy, there is an algorithm that decides whether or not \mathbf{a} is in the core of a public goods economy. The computational bottleneck in this algorithm is the solution of $O(n)$ convex programming problems on the utility function of the economy.*

Hence we essentially have the first polynomial-time tester for coalitional stability in an unrestricted public goods economy, up to the implementation of the necessary convex programming oracle. It is fairly natural in our economic metaphor to assume that convex programming should be tractable: it corresponds to the negotiation process of a group of agents trying to determine how well they can maximize a joint utility function as a group. If even this is impossible, then it is essentially hopeless to efficiently test core membership. One cannot even test the more basic property of Pareto efficiency, a necessary step towards testing core membership, without assuming some computational power along these lines.

Even so, if one does not wish to introduce such assumptions, Theorem 3 implies that several broad special cases of membership testing have efficient algorithms. The most obvious of these is when the utility function and its derivative can be described by a “nice” analytic function on which the standard derivative-based method for exact convex optimization goes through. Less obviously, if convex optimization really is hard for the given utility function (or the utility function of the economy is unknown), one can employ modern ε -approximate convex optimization solvers, which treat the utility function as a black box that can be queried, to solve certain natural ε -approximate relaxations of the core membership testing problem. We discuss this point in the conclusion of the paper, since it is easier to be specific here once the economic model is familiar.

We consider it somewhat surprising that these algorithmic results are possible, given a general dearth of positive results in the area. Moreover, the approach taken by the algorithm is fairly intuitive and seems to plausibly reflect practical behavior. Starting with the grand coalition, we show (via Theorem 1) that we can either determine that the current coalition has a deviation, or we can identify a “least-valuable player” who is formally the least likely agent to participate in a deviation. We then kill this agent and repeat the analysis on the survivors. After n rounds, we have either killed every agent (and thus determined that the given point is coalitionally stable), or we have explicitly found a surviving coalition with a deviation. It is quite reasonable to imagine that a practical search for a deviating coalition might employ a “greedy” method of iteratively killing the agent who seems to be least pulling their weight at the current agreement; an insight of Theorem 3 is that this search heuristic is in fact thorough and will provably produce the right answer.

1.4 Comparison with Prior Work

Elliot and Golub [12] recently studied the Lindahl equilibria in public goods economies, with a focus on characterizing the set of solutions rather than algorithmically computing/testing them. More specifically, they frame the typical model of public goods economies in the language of networks, and use this to equate the eigenvectors of the “benefits network” with

the Lindahl equilibria of the economy. A less general version of this networks interpretation was implicitly used in several other papers concerning Nash equilibria of public goods economies, for example [3, 4, 1, 2]. In this paper, we will adopt the more general networks-based phrasing of public goods economies used by Elliot and Golub, and we will rely on this insight in a critical way to prove our main results.

Per the discussion above, there has been lots of prior work on the algorithmic properties of the core, largely intended to confirm/refute the bounded rationality argument in some economic model. Three questions are commonly studied:

- The *membership testing* problem (discussed above): is a given outcome in the core of the game?
- The *emptiness testing* problem: is the core empty?
- The *member finding* problem: output any solution in the core of the game (if nonempty).

We remark that the latter two problems are already closed in public goods economies: Elliot and Golub [12] show that the core is never empty except in certain degenerate cases, and it can be seen from the model below that the member finding problem is essentially identical to the general problem of convex optimization (which is well beyond the scope of this economically-minded research program). Hence, this work is entirely focused on the membership testing problem.

In order to frame these three questions as proper computational problems, past work has commonly defined a “compressed” cooperative game that allows the payoffs achievable by all 2^n possible coalitions to be expressed on only $\text{poly}(n)$ input bits. For example, in a seminal paper by Deng and Papadimitriou introducing this line of research [9], the authors studied *graphical games* in which weighted edges are placed between agents and the value attainable by a coalition is equal to the total weight contained in its induced subgraph. Upper and lower bounds are often obtained for these problems by exploiting particular features of the compression scheme. By contrast, our goal is to assume as little structure for the problem as possible (since our main results are upper bounds, this is the more general approach). Thus, we allow the economy to be governed by an arbitrarily complex utility function, which does not need to have a succinct representation, or even any algorithmic representation at all. Instead, we allow ourselves black-box constant-time query access to the utility function, which acts as an oracle and thus may have arbitrary complexity. The goal in this substantial generalization is to ensure that our results reveal structure of the core itself, rather than the nature of an assumed compression.

2 The Model and Basic Definitions

2.1 Notation Conventions

Given vectors $\mathbf{a}, \mathbf{a}' \in \mathbb{R}^n$, we will use the following (partial) ordering operations:

- $\mathbf{a} \geq \mathbf{a}'$ means that $a_i \geq a'_i$ for all $1 \leq i \leq n$,
- $\mathbf{a} > \mathbf{a}'$ means that $a_i > a'_i$ for all $i \leq i \leq n$, and
- $\mathbf{a} \succeq \mathbf{a}'$ means that $a_i \geq a'_i$ for all $1 \leq i \leq n$, and $a_j > a'_j$ for some $1 \leq j \leq n$.

Given a subset $C \subseteq \{1, \dots, n\}$ and a vector $\mathbf{v} \in \mathbb{R}^n$, we write \mathbf{v}_C to denote the restriction of \mathbf{v} to the indices in C ; that is, \mathbf{v}_C is the $|C|$ length vector built by deleting the entry v_i from \mathbf{v} for each $i \notin C$.

We use $\mathbf{0}, \mathbf{1}$ as shorthands for the vectors $\langle 0, \dots, 0 \rangle, \langle 1, \dots, 1 \rangle$ respectively.

2.2 Economic Model

We adopt the terminology of Elliot and Golub [12] when possible. The salient pieces of our economy are defined as follows:

- The set of agents in the economy is given by $N = [n] = \{1, \dots, n\}$. A nonempty subset of agents in the economy $C \subseteq N$ is called a *coalition*. The coalition $C = N$ is called the *grand coalition*.
- Each agent i chooses an *action* a_i , which can be any real number in the interval $[0, 1]$. An *outcome* or *point* is a vector $\mathbf{a} \in \mathbb{R}^n$ built by concatenating the actions of all agents.
- There is a continuous *utility function* $\mathbf{u} : [0, 1]^n \rightarrow [0, 1]^n$, which maps outcomes to a level of “utility” for each agent. In particular, agent i prefers outcome \mathbf{a} to outcome \mathbf{a}' iff $u_i(\mathbf{a}) > u_i(\mathbf{a}')$. The utility function has the following two properties:
 - *Positive Externalities*: whenever $\mathbf{a} \succeq \mathbf{a}'$ with $a_i = a'_i$, we have $u_i(\mathbf{a}) > u_i(\mathbf{a}')$. This assumption is what places us in the setting of public goods economies; intuitively, it states that an agent gains utility when other agents increase their production of public goods.
 - *Convex Preferences*: we assume that \mathbf{u} is concave.³ That is, for any outcomes \mathbf{a}, \mathbf{a}' and any $\lambda \in [0, 1]$, we have $\mathbf{u}(\lambda\mathbf{a} + (1 - \lambda)\mathbf{a}') \geq \lambda\mathbf{u}(\mathbf{a}) + (1 - \lambda)\mathbf{u}(\mathbf{a}')$. This standard assumption corresponds to the economic principle of diminishing marginal returns.

2.3 Game Theory Definitions

We recap some well-known definitions from the game theory literature.

► **Definition 4** (Pareto Efficiency). An outcome \mathbf{a} is a *Pareto Improvement* on another outcome \mathbf{a}' if $\mathbf{u}(\mathbf{a}) \succeq \mathbf{u}(\mathbf{a}')$. An outcome \mathbf{a} is *Pareto Efficient* if there is no Pareto improvement on \mathbf{a} . The set of Pareto efficient outcomes is called the *Pareto Frontier*.

The main solution concept that will be discussed in this paper is *the core*:

► **Definition 5** (Deviation). Given an outcome \mathbf{a} , an outcome \mathbf{a}' is a *deviation* from \mathbf{a} for a coalition C if $\mathbf{a}'_{N \setminus C} = \mathbf{0}$ and $\mathbf{u}_C(\mathbf{a}') > \mathbf{u}_C(\mathbf{a})$.

► **Definition 6** (The Core). An outcome \mathbf{a} is in the *core* of the economy if no coalition has a deviation from \mathbf{a} (equivalently, \mathbf{a} is *coalitionally stable*).

The next definition that will be useful in our proofs is the *projected economy*:

► **Definition 7**. Given an economy described by agents N and a utility function \mathbf{u} , the *projected economy* for a coalition C is the economy described by agents N and utility function $\mathbf{u}^C(\mathbf{a}_C)$, where

$$\mathbf{u}^C(\mathbf{a}_C) := \mathbf{u}_C(\mathbf{a}_C \cdot \mathbf{0}_{N \setminus C}).$$

In other words, the new $|C|$ -dimensional utility function \mathbf{u}^C is obtained by fixing the actions of $N \setminus C$ at 0, allowing any action for C , and then using the old utility function \mathbf{u} to determine the utilities for C in the natural way. We suppress the superscript \mathbf{u}^C when clear from context.

³ Confusingly, when \mathbf{u} is mathematically concave, one says that preferences are “economically convex” – hence, *convex preferences*.

► **Definition 8.** The *dominated set* of \mathbf{a} , denoted $\mathbf{D}_{\mathbf{a}}$, is defined as:

$$\mathbf{D}_{\mathbf{a}} := \{\mathbf{a}' \mid \mathbf{u}(\mathbf{a}) \geq \mathbf{u}(\mathbf{a}')\}.$$

In other words, $\mathbf{D}_{\mathbf{a}}$ is the set of points that no agent prefers to \mathbf{a} . Note that this is an unusually weak definition of dominance, in the sense that (for example) $\mathbf{D}_{\mathbf{a}}$ contains \mathbf{a} itself.

3 A Topological Characterization of the Core

Our goal in this section is to prove the following structural theorem:

► **Theorem 9.** *Let \mathbf{a} be an outcome in a public goods economy. Then \mathbf{a} is in the core if and only if it is Pareto efficient, individually rational, and $\mathbf{D}_{\mathbf{a}}$ is path connected.*

The vast majority of the technical depth of this theorem is tied up in the implication

$$\mathbf{a} \text{ is in the core} \longrightarrow \mathbf{D}_{\mathbf{a}} \text{ is path connected.}$$

The remainder of this forwards implication (\mathbf{a} is in the core \rightarrow neither the grand coalition nor any singleton coalition has a deviation from \mathbf{a}) is extremely straightforward: \mathbf{a} is individually rational iff each agent i prefers it to the outcome they can guarantee acting alone, which coincides with the notion that the singleton coalition $\{i\}$ has no deviation from \mathbf{a} . In our model, Pareto efficiency coincides with the notion that the grand coalition N has no deviation from \mathbf{a} :

► **Claim 10.** *Let \mathbf{a} be an outcome. If there is an outcome \mathbf{a}' satisfying $\mathbf{a} \succeq \mathbf{a}'$ ($\mathbf{a} \preceq \mathbf{a}'$) and $\mathbf{u}(\mathbf{a}) \succeq \mathbf{u}(\mathbf{a}')$, then there is an outcome \mathbf{a}'' satisfying $\mathbf{a} > \mathbf{a}''$ ($\mathbf{a} < \mathbf{a}''$) and $\mathbf{u}(\mathbf{a}) > \mathbf{u}(\mathbf{a}'')$.*

Proof. We will prove the claim for the case $\mathbf{a} \succeq \mathbf{a}'$; the case $\mathbf{a} \preceq \mathbf{a}'$ follows from a symmetric argument.

Choose an agent i for whom $u_i(a) > u_i(a')$, and then slightly increase a_i . Since \mathbf{u} is continuous, if we increase a_i by a sufficiently small amount then we still have $u_i(a) > u_i(a')$. Additionally, by positive externalities we then have $\mathbf{u}(\mathbf{a}) > \mathbf{u}(\mathbf{a}')$. We can then slightly increase the actions of all agents, such that $\mathbf{a} > \mathbf{a}'$, but with sufficiently small increases we do not destroy the property that $\mathbf{u}(\mathbf{a}) > \mathbf{u}(\mathbf{a}')$. ◀

In this section, we will first give a complete proof of the (easier) backwards implication of Theorem 9, and then we sketch the proof of the forwards implication. Due to space constraints, a full proof of the forwards implication can be found in the full version of this paper.

3.1 Backwards Implication of Theorem 9

First:

► **Lemma 11.** *If \mathbf{a} is Pareto efficient, then every deviation \mathbf{a}' from \mathbf{a} satisfies $\mathbf{a}' \preceq \mathbf{a}$.*

Proof. Let I be the set of agents i for which $a'_i > a_i$, and suppose towards a contradiction that I is nonempty.

Consider the point \mathbf{a}'' defined such that $\mathbf{a}''_{N \setminus I} := \mathbf{a}_{N \setminus I}$ and $\mathbf{a}''_I := \mathbf{a}'_I$. We then have $\mathbf{u}_{N \setminus I}(\mathbf{a}'') > \mathbf{u}_{N \setminus I}(\mathbf{a})$ by positive externalities, since these points differ only in that the (nonempty) coalition I has increased their actions. We also have $\mathbf{u}_I(\mathbf{a}'') \geq \mathbf{u}_I(\mathbf{a}') > \mathbf{u}_I(\mathbf{a})$, where the first inequality follows from positive externalities (since these points differ only in

that the coalition $N \setminus I$ has weakly increased their action), and the second follows from the fact that \mathbf{a}' is a deviation from \mathbf{a} for a coalition C with $I \subseteq C$ (since $\mathbf{a}_I > \mathbf{a}_I$).

We thus have $\mathbf{u}(\mathbf{a}'') > \mathbf{u}(\mathbf{a})$, which contradicts the fact that \mathbf{a} is Pareto efficient. Thus I is empty and the lemma follows. ◀

Second:

► **Lemma 12.** *Suppose there is a path $P \subset \mathbb{R}_{\geq 0}^n$ with endpoints \mathbf{x}, \mathbf{y} such that for any $\mathbf{p} \in P$ we have $\mathbf{u}(\mathbf{p}) \leq \mathbf{u}(\mathbf{a})$. If \mathbf{a}' is a deviation from \mathbf{a} for some coalition C satisfying $\mathbf{a}' \preceq \mathbf{x}$, then \mathbf{a}' also must satisfy $\mathbf{a}' \preceq \mathbf{y}$.*

Proof. We walk along P from \mathbf{x} towards \mathbf{y} until we find the first point \mathbf{p} with $p_i = a'_i \neq 0$ for some i . If we reach \mathbf{y} before we find any such point \mathbf{p} , it follows that $a'_i = 0$ or $a'_i < x_i$ for all i , and so $\mathbf{a}' \preceq \mathbf{x}$, as claimed. Otherwise, we find such a point \mathbf{p} , and we argue towards a contradiction.

We have $\mathbf{p} \neq \mathbf{a}'$, since $\mathbf{u}_C(\mathbf{a}') \succeq \mathbf{u}_C(\mathbf{a})$ but $\mathbf{u}_C(\mathbf{p}) \leq \mathbf{u}_C(\mathbf{a})$. By construction we then have $\mathbf{p} \succeq \mathbf{a}'$. Since $p_i = a'_i$, by positive externalities we then have $u_i(p) > u_i(a')$. Since $a'_i \neq 0$ we have $i \in C$, and since \mathbf{a}' is a deviation for C , this implies $u_i(a') > u_i(a)$. We then have $u_i(p) > u_i(a)$, which contradicts the assumption that $\mathbf{u}(\mathbf{p}) \leq \mathbf{u}(\mathbf{a})$. Therefore no such point \mathbf{p} may be found. ◀

We can now show:

Proof of Theorem 9, Backwards Implication. Assume that \mathbf{a} is robust to deviations by the grand coalition or any singleton coalition, and that $\mathbf{D}_{\mathbf{a}}$ is path connected. Our goal is now to show that \mathbf{a} is in the core.

By Claim 10, the property that the grand coalition has no deviation from \mathbf{a} implies that \mathbf{a} is Pareto efficient. Thus, by Lemma 11 any deviation \mathbf{a}' from \mathbf{a} satisfies $\mathbf{a}' \preceq \mathbf{a}$. Since no singleton coalition has a deviation from \mathbf{a} we have $\mathbf{0} \in \mathbf{D}_{\mathbf{a}}$, and since $\mathbf{D}_{\mathbf{a}}$ is path connected there is a path contained in $\mathbf{D}_{\mathbf{a}}$ with endpoints \mathbf{a}, \mathbf{a} . Thus, by Lemma 12, we further have that a deviation \mathbf{a}' must satisfy $\mathbf{a}' \preceq \mathbf{0}$. Since no such point exists, it follows that no deviations from \mathbf{a} exist, and so \mathbf{a} is a core outcome. ◀

3.2 Sketch of Forwards Implication of Theorem 9

We will denote by $\mathbf{d}_{\mathbf{v}}\mathbf{u}(\mathbf{a})$ the one-sided directional derivative of \mathbf{u} at \mathbf{a} in the direction \mathbf{v} . In other words:

$$\mathbf{d}_{\mathbf{v}}\mathbf{u}(\mathbf{a}) := \lim_{\lambda \rightarrow 0^+} \frac{\mathbf{u}(\mathbf{a} + \lambda\mathbf{v}) - \mathbf{u}(\mathbf{a})}{\lambda}.$$

A nontrivial but standard fact from analysis is that, since \mathbf{u} is concave and well-defined everywhere, this limit is well-defined for all \mathbf{a} , except when excluded by a boundary condition (e.g. if $v_i < 0$ but $a_i = 0$ for some agent i) – see [21].

Our key lemma is:

► **Lemma 13.** *At any outcome $\mathbf{0} < \mathbf{a} < \mathbf{1}$, exactly one of the following three conditions holds:*

1. *There exist directions $\mathbf{v}_{up} > \mathbf{0}, \mathbf{v}_{down} < \mathbf{0}$ such that $\mathbf{d}_{\mathbf{v}_{up}}\mathbf{u}(\mathbf{a}) > \mathbf{0}$ and $\mathbf{d}_{\mathbf{v}_{down}}\mathbf{u}(\mathbf{a}) < \mathbf{0}$,*
2. *There exist directions $\mathbf{v}_{up} > \mathbf{0}, \mathbf{v}_{down} < \mathbf{0}$ such that $\mathbf{d}_{\mathbf{v}_{up}}\mathbf{u}(\mathbf{a}) \leq \mathbf{0}$ and $\mathbf{d}_{\mathbf{v}_{down}}\mathbf{u}(\mathbf{a}) \leq \mathbf{0}$,*
or
3. *There exist directions $\mathbf{v}_{up} > \mathbf{0}, \mathbf{v}_{down} < \mathbf{0}$ such that $\mathbf{d}_{\mathbf{v}_{up}}\mathbf{u}(\mathbf{a}) < \mathbf{0}$ and $\mathbf{d}_{\mathbf{v}_{down}}\mathbf{u}(\mathbf{a}) > \mathbf{0}$.*

The three categories of Lemma 13 carry a useful geometric intuition. Specifically:

► **Lemma 14.** *The points in the second category of Lemma 13 are precisely the Pareto Frontier.*

The proofs of these two lemmas are quite technical, and can be found in the full version of this paper. With these in mind, we define

► **Definition 15.** We will say that a point in the first category of Lemma 13 is *below the Pareto Frontier*, and a point in the third category of Lemma 13 is *above the Pareto Frontier* (and by Lemma 14, the second category of points in Lemma 13 are on the Pareto Frontier).

The geometric intuition behind this definition is that, starting from a point \mathbf{x} in the first category, one can continuously follow the gradient \mathbf{v}_{up} to eventually obtain a Pareto efficient Pareto improvement $\mathbf{x}' > \mathbf{x}$ (we do not prove this fact formally; it is perhaps useful intuition but not essential to our main results). Similarly, starting from \mathbf{x} in the third category, we can continuously follow the gradient \mathbf{v}_{down} to obtain a Pareto efficient Pareto improvement $\mathbf{x}' < \mathbf{x}$.

We then show:

Proof Sketch of Theorem 9, Forwards Implication (Proof in full version). Suppose \mathbf{a} is a core outcome, and our goal is to show path connectedness of $\mathbf{D}_{\mathbf{a}}$. First, we note that $\mathbf{0} \in \mathbf{D}_{\mathbf{a}}$, since otherwise a singleton coalition can deviate from \mathbf{a} (and \mathbf{a} is in the core, so no such deviation is possible). To show path connectedness of $\mathbf{D}_{\mathbf{a}}$, we consider an arbitrary point $\mathbf{x} \in \mathbf{D}_{\mathbf{a}}$ and construct a path in $\mathbf{D}_{\mathbf{a}}$ from \mathbf{x} to $\mathbf{0}$, thus implying that any two such points $\mathbf{x}, \mathbf{x}' \in \mathbf{D}_{\mathbf{a}}$ have a connecting path in $\mathbf{D}_{\mathbf{a}}$ via $\mathbf{0}$.

We show the existence of the $\mathbf{x} \rightsquigarrow \mathbf{0}$ path with a careful repeated application of Lemma 13. Informally speaking, we progressively slide $\mathbf{x} > \mathbf{0}$ a little bit closer to $\mathbf{0}$ while maintaining the property $\mathbf{x} \in \mathbf{D}_{\mathbf{a}}$. If we ever hit $x_i = 0$ for some agent i , then we restrict our attention to the projected economy disclosing agent i and continue. If we eventually exclude all agents in this manner, then we have $\mathbf{x} = \mathbf{0}$ and the process is complete. Otherwise, suppose towards a contradiction that at some \mathbf{x} , we cannot slide \mathbf{x} any closer to $\mathbf{0}$ while maintaining $\mathbf{x} \in \mathbf{D}_{\mathbf{a}}$. We make two observations here: (1) \mathbf{x} must be above the Pareto frontier (else we could slide \mathbf{x} in the appropriate direction \mathbf{v}_{down}) and so it belongs to the third; and (2) for all agents i still being considered, we have $u_i(\mathbf{x}) = u_i(\mathbf{a})$ (else, by the positive externalities assumption, we can unilaterally decrease the action of agent i without destroying $\mathbf{x} \in \mathbf{D}_{\mathbf{a}}$). Hence, by moving \mathbf{x} slightly in the direction \mathbf{v}_{down} (which *improves* the utility of all agents being considered), we have $u_i(\mathbf{x}) > u_i(\mathbf{a})$ for all agents being considered, and so the new \mathbf{x} is a deviation from \mathbf{a} . Since we have assumed that \mathbf{a} is a core outcome, this is a contradiction, and so the process of sliding \mathbf{x} towards $\mathbf{0}$ can never get stuck in this way. ◀

3.3 Connection to Lindahl Equilibria

Before proceeding towards our algorithm, we take a brief detour in this subsection to observe an interesting implication of Theorem 9 that helps illustrate its broader appeal. Elliot and Golub [12] show the following result:

► **Theorem 16** ([12]). *The Lindahl equilibria of a public goods economy with a differentiable utility function are precisely the outcomes \mathbf{a} for which $\mathbf{d}_{\mathbf{a}}\mathbf{u}(\mathbf{a}) = \mathbf{0}$.*

They phrase this theorem in different language related to the “benefits network” of the economy, but this formulation will suit our purposes better. We refer the reader to their paper for a more in-depth discussion of the economic role of the Lindahl equilibria.

45:10 Testing Core Membership in Public Goods Economies

Combining Theorem 16 with our machinery for Theorem 9, we obtain:

► **Theorem 17.** *In a public goods economy with a differentiable utility function, the Lindahl equilibria are precisely the core outcomes \mathbf{a} whose membership can be certified by examining only local information at \mathbf{a} .*

The proof of this theorem will use Theorem 16 as a black box, and so we will not actually need to appeal to the formal definition of the Lindahl equilibria in its proof.

Proof. If \mathbf{a} is a Lindahl equilibrium, then by Theorem 16 we have $\mathbf{d}_{\mathbf{a}} \mathbf{u}(\mathbf{a}) = \mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) = \mathbf{0}$ (the first equality comes from the assumption that \mathbf{u} is differentiable). We claim that any \mathbf{a} satisfying $\mathbf{d}_{\mathbf{a}} \mathbf{u}(\mathbf{a}) = \mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) = \mathbf{0}$ is in the core, and thus its core membership can be verified by examining only these local derivatives. First, note that \mathbf{a} is Pareto efficient by Lemma 14. Therefore, by Lemma 11, any possible deviation \mathbf{a}' from \mathbf{a} satisfies $\mathbf{a}' \preceq \mathbf{a}$. Now let P be the line segment from $\mathbf{0}$ to \mathbf{a} . By the assumption of concavity and the fact that $\mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) = \mathbf{0}$, we have $\mathbf{u}(\mathbf{p}) \leq \mathbf{u}(\mathbf{a})$ for all $\mathbf{p} \in P$. Thus, by Lemma 12, we have $\mathbf{a}' \preceq \mathbf{0}$ and so \mathbf{a}' cannot exist and \mathbf{a} is a core outcome.

Now suppose that \mathbf{a} is not a Lindahl equilibrium, and so $\mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) \neq \mathbf{0}$. If we have $\mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) \leq \mathbf{0}$, then we have $\mathbf{d}_{\mathbf{a}} \mathbf{u}(\mathbf{a}) \geq \mathbf{0}$ (by differentiability) which implies that \mathbf{a} is not Pareto efficient, and hence is not a core outcome. On the other hand, suppose we have $\mathbf{d}_{-\mathbf{a}} \mathbf{u}(\mathbf{a}) \not\leq \mathbf{0}$. In this case, it is impossible to distinguish \mathbf{u} from the utility function \mathbf{u}' that is affine-linear everywhere and agrees with \mathbf{u} at \mathbf{a} using solely local information. Note that \mathbf{a} is *not* in the core of the economy defined by \mathbf{u}' , since we have $u_i(\mathbf{0}) > u_i(\mathbf{a})$ for whichever agent i satisfies $d_{-a} u_i(\mathbf{a}) > 0$. Thus, if \mathbf{a} is in fact in the core of the economy defined by \mathbf{u} , we will need to inspect non-local information about the economy to differentiate \mathbf{u} from \mathbf{u}' . ◀

We note that it is possible to prove Theorem 17 as a corollary directly from Theorem 9, but this proof using the underlying machinery is simpler.

4 Algorithm for Testing Core Membership

Our main algorithmic result is:

► **Theorem 18.** *Given an outcome \mathbf{a} in a public goods economy, there is an algorithm (in the real-RAM model) that decides whether or not \mathbf{a} is in the core by solving $O(n)$ convex optimization problems and using $O(n)$ additional computation time.*

The algorithm is fairly straightforward. We maintain an “active coalition” C_A throughout, as well as a proof that any agent $i \notin C_A$ must play action $a'_i = 0$ in any deviation \mathbf{a}' from \mathbf{a} . It is thus safe to assume that any deviating coalition C satisfies $C \subseteq C_A$. Initially $C_A \leftarrow N$, so this invariant is trivially satisfied. After each round, we either find a deviation for C_A from \mathbf{a} , or we remove one new agent from C_A . Thus, if we make it n rounds without finding a deviation, then we have $C_A = \emptyset$ and so no deviation from \mathbf{a} is possible.

4.1 Preprocessing: Confirm Pareto Efficiency of \mathbf{a}

Before starting the main algorithm, we run the following two programs, with the purpose of testing whether or not \mathbf{a} is Pareto efficient.

► **Program 1.** Choose \mathbf{v} to maximize $\min_i d_v u_i(\mathbf{a})$

$$\text{Subj. to } \mathbf{v} \geq \mathbf{0}, \sum_i v_i = 1$$

- **Program 2.** Choose \mathbf{v} to maximize $\min_i d_v u_i(a)$
 Subj. to $\mathbf{v} \leq \mathbf{0}, \sum_i v_i = 1$

Note that the concavity of the optimized function $f(v) := \min_i d_v u_i(a)$ is immediate from the concavity of \mathbf{u} . By Lemma 13, we may immediately conclude that \mathbf{a} is not Pareto efficient (and thus not in the core) iff either of these programs optimizes at a point \mathbf{v}^* satisfying $f(v^*) > 0$. Otherwise, we proceed with the knowledge that \mathbf{a} is Pareto efficient. A key advantage of this is that, by Lemma 11, we may now restrict our search for a deviation \mathbf{a}' to the bounded box $\mathbf{0} \leq \mathbf{a}' \leq \mathbf{a}$. This opens up the ability to use convex programming algorithms, which typically require bounded domains, in the remainder of the algorithm.⁴

4.2 Main Loop: Shrinking C_A

Each of the n rounds of the algorithm consists of three steps. First, we restrict our attention to the projected economy for the coalition C_A . Second, we run the following program:

- **Program 3.** Choose \mathbf{x} to maximize $\min_i u_i(x) - u_i(a)$
 Subj. to $\mathbf{0} \leq \mathbf{x} \leq \mathbf{a}_{C_A}$

Let \mathbf{x}^* be a maximizing point of Program 3. We have:

- **Lemma 19.** *Either $\mathbf{u}^{C_A}(\mathbf{x}^*) > \mathbf{u}_{C_A}(\mathbf{a})$ or $\mathbf{u}^{C_A}(\mathbf{x}^*) \leq \mathbf{u}_{C_A}(\mathbf{a})$, and \mathbf{x}^* is Pareto efficient.*⁵

Proof. First we argue Pareto efficiency. If \mathbf{x}' is a Pareto improvement on \mathbf{x}^* , then by Claim 10 there is another point \mathbf{x}'' with $\mathbf{u}(\mathbf{x}'') > \mathbf{u}(\mathbf{x}^*)$. This \mathbf{x}'' would be a superior maximizing point for Program 3, so there can be no Pareto improvement on \mathbf{x}^* .

Next, let $i := \arg \max_i u_i(x^*) - u_i(a)$ and $j := \arg \min_j u_j(x^*) - u_j(a)$. If $u_i(x^*) - u_i(a) > u_j(x^*) - u_j(a)$, then (by the same argument used in Claim 10) we can again obtain a superior maximizing point \mathbf{x}^{**} by slightly increasing the action of agent i from \mathbf{x}^* . Thus we have $u_i(x^*) - u_i(a) = u_j(x^*) - u_j(a)$, and it follows that either $\mathbf{u}(\mathbf{x}^*) > \mathbf{u}(\mathbf{a})$ or $\mathbf{u}(\mathbf{x}^*) \leq \mathbf{u}(\mathbf{a})$. ◀

In the former case where $\mathbf{u}(\mathbf{x}^*) > \mathbf{u}(\mathbf{a})$, it follows that \mathbf{x}^* is a deviation from \mathbf{a} for the coalition C_A , so we may halt the algorithm. Otherwise, we have $\mathbf{u}(\mathbf{x}^*) \leq \mathbf{u}(\mathbf{a})$. We then observe:

- **Lemma 20.** *If $\mathbf{u}^{C_A}(\mathbf{x}^*) \leq \mathbf{u}_{C_A}(\mathbf{a})$, then in the full (non-projected) economy, any deviation \mathbf{a}' from \mathbf{a} for a coalition $C \subseteq C_A$ satisfies $\mathbf{a}'_{C_A} \preceq \mathbf{x}^*$.*

Proof. The deviation \mathbf{a}' satisfies $\mathbf{a}'_{N \setminus C} = \mathbf{0}$ and $\mathbf{u}_C(\mathbf{a}') > \mathbf{u}_C(\mathbf{a}) \geq \mathbf{u}_C(\mathbf{x}^*)$. It follows that \mathbf{a}'_C is also a deviation for C from \mathbf{x}^* in the projected economy for C_A . The claim is then immediate from Lemma 11. ◀

One step remains. We run:

- **Program 4.** Choose \mathbf{v} to maximize $\min_i d_v u_i(x^*)$
 Subj. to $\mathbf{v} \leq \mathbf{0}, \sum_i v_i = 1$

⁴ This detail is precisely why we use Programs 1 and 2 to check the Pareto efficiency of \mathbf{a} , rather than the ostensibly simpler method of searching for \mathbf{x}^* that maximizes $\min_i u_i(x^*) - u_i(a)$: the latter method requires a search for \mathbf{x}^* over an unbounded search space, which rules out many popular methods of convex optimization that we wish to keep available.

⁵ Note that these statements hold specifically in the projected economy for C_A .

Algorithm 1: Testing Core Membership of \mathbf{a} .

```

1 Let  $\mathbf{v}_1^* \leftarrow$  output of Program 1;
2 Let  $\mathbf{v}_2^* \leftarrow$  output of Program 2;
3 if  $\mathbf{d}_{\mathbf{v}_1^*} \mathbf{u}(\mathbf{a}) > \mathbf{0}$  or  $\mathbf{d}_{\mathbf{v}_2^*} \mathbf{u}(\mathbf{a}) > \mathbf{0}$  then
4   | return “ $\mathbf{a}$  is not in the core”;
5 end
6  $C_A \leftarrow N$ ;
7 while  $C_A \neq \emptyset$  do
8   |  $\mathbf{x}^* \leftarrow$  output of Program 3;
9   | if  $\mathbf{u}^{C_A}(\mathbf{x}^*) > \mathbf{u}_{C_A}(\mathbf{a})$  then
10    | | return “ $\mathbf{a}$  is not in the core”;
11    | end
12    |  $\mathbf{v}^* \leftarrow$  output of Program 4;
13    |  $C_A \leftarrow C_A \setminus \{\arg \min_{i \in C_A} x_i^*/v_i^*\}$ ;
14  end
15 return “ $\mathbf{a}$  is in the core”;

```

Let \mathbf{v}^* be a point that maximizes Program 4. We have

► **Lemma 21.**

$$\mathbf{d}_{\mathbf{v}^*} \mathbf{u}(\mathbf{x}^*) \leq \mathbf{0}.$$

The proof is very similar to the proof of Lemma 19, so we omit it for now. We then finally have:

► **Lemma 22.** *Let $i := \arg \min_i x_i^*/v_i^*$. Then any deviation \mathbf{a}' from \mathbf{a} has $a'_i = 0$.*

Proof. By Lemma 20, we have $\mathbf{a}'_{C_A} \preceq \mathbf{x}^*$. Let P be the line segment starting at \mathbf{x}^* , extending in the direction \mathbf{v}^* until a point \mathbf{p}^* is reached where $p_i^* = 0$ for some agent i ; note that this will specifically be $i = \arg \min_i x_i^*/v_i^*$. By concavity, all $\mathbf{p} \in P$ satisfies $\mathbf{u}(\mathbf{p}) \leq \mathbf{u}(\mathbf{x}^*) \leq \mathbf{u}(\mathbf{a})$. Noting once again that \mathbf{a}'_{C_A} is a deviation for C from \mathbf{x}^* in the projected economy for C_A , it follows from Lemma 12 that $\mathbf{a}'_{C_A} \preceq \mathbf{p}^*$, and so $a'_i = 0$. ◀

With this in mind, the final step in the loop is to delete i from C_A and repeat. After n repetitions, we have $C_A = \emptyset$, so we may halt the algorithm and report that \mathbf{a} is in the core.

4.3 Algorithm Pseudocode

To recap the algorithm, which has been interspersed with proofs of correctness above, we give full pseudocode here.

4.4 Conclusion

Our algorithm implies that core membership testing is efficient under any utility function that admits quick solving of convex programs as described above. However, it may still be desirable to test core membership as best as possible when the underlying utility function is either unknown or badly behaved and so exact convex optimization is impossible. Our algorithm can indeed be adapted to this effect, with a few significant points of caution, by substituting in modern approximate optimization algorithms. Due to space constraints, we defer a discussion of this point to the full version of the paper.

References

- 1 N. Allouch. The cost of segregation in social networks. *Queen Mary Working Paper*, 2013.
- 2 N. Allouch. On the private provision of public goods on networks. *Journal of Economic Theory*, pages 527–552, 2015.
- 3 C. Ballester, A. Calvó-Armengol, and Y. Zenou. Who’s who in networks. wanted: The key player. *Econometrica*, 74:1403–1417, 2006.
- 4 Y. Bramouille and R. Kranton. Public goods in networks. *Journal of Economic Theory*, 135:478–494, 2007.
- 5 Parkash Chander and Henry Tulkens. The core of an economy with multilateral environmental externalities. *International Journal of Game Theory*, 26:379–401, 1997.
- 6 V. Conitzer and T. Sandholm. Computing shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *Proc. National Conference on Artificial Intelligence (AAAI)*, pages 219–225, 2004.
- 7 V. Conitzer and T. Sandholm. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence*, pages 607–619, 2006.
- 8 Imma Curriel. *Cooperative Game Theory and Applications*. Springer US, 1 edition, 1997.
- 9 X. Deng and C. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19:257, 1994.
- 10 F. Y. Edgeworth. Mathematical psychics: An essay on the mathematics to the moral sciences. *Reprinted in Diamond, M.A.*, 1881.
- 11 E. Elkind, L. A. Goldberg, P. W. Goldberg, and M. Wooldridge. A tractable and expressive class of marginal contribution nets and its applications. *Math. Log. Q.*, 2009.
- 12 Matthew Elliot and Benjamin Golub. A network approach to public goods. In *Proc. 14th Electronic Commerce (EC)*, pages 377–378, 2013.
- 13 U. Faigle, S. Fekete, W. Hochstattler, and W. Kern. On the complexity of testing membership in the core of min-cost spanning trees. *Internat. J. Game Theory*, 26:361–366, 1997.
- 14 D. K. Foley. Lindahl’s solution and the core of an economy with public goods. *Econometrica*, 38:66–72, 1970.
- 15 D. B. Gillies. Solutions to general non-zero-sum games. *Contributions to the Theory of Games IV*, pages 47–85, 1959.
- 16 M. Goemans and M. Skutella. Cooperative facility location games. In *Symposium on Discrete Algorithms (SODA)*, pages 76–85, 2000.
- 17 Benjamin Golub. Personal correspondence, 2012.
- 18 G. Greco, E. Malizia, L. Palopoli, and F. Scarcello. On the complexity of the core over coalition structures. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- 19 S. Yeung and Y. Shoham. Marginal contribution nets: A compact representation scheme for coalitional games. In *Electronic Commerce (EC)*, pages 193–202, 2005.
- 20 S. Kintali, L. Poplawski, R. Rajaraman, R. Sundaram, and S. Teng. Reducibility among fractional stability problems. In *Proc. 50th FOCS*, pages 283–292, 2009.
- 21 Jonathan Lewin. *An interactive introduction to mathematical analysis*. Cambridge University Press, 2003.
- 22 Y. Li and V. Conitzer. Complexity of stability-based solution concepts in multi-issue and mc-net cooperative games. In *Proc. 2014 international conference on Autonomous agents and multi-agent systems (AAMAS)*. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- 23 Bezalel Peleg and Peter Sudhölter. *Introduction to the Theory of Cooperative Games*. Springer, 2 edition, 10 2007.

45:14 Testing Core Membership in Public Goods Economies

- 24 Paul Samuelson. The pure theory of public expenditure. *Review of Economics and Statistics*, 37(4):350–356, 1954.
- 25 H. Scarf. The core of an n person game. *Econometrica*, 1967.
- 26 S. Sung and D. Dimitrov. On core membership testing for hedonic coalition formation games. *Operations Research Letters*, 35:155–158, 2007.