

Non-Myopic Negotiators See What’s Best

Yair Zick
Carnegie Mellon University
yairzick@cs.cmu.edu

Yoram Bachrach and Ian A. Kash and Peter Key
Microsoft Research
yobach,iankash,Peter.Key@microsoft.com

Abstract

We consider revenue negotiation problems in iterative settings. In our model, a group of agents has some initial resources, used in order to generate revenue. Agents must agree on some way of dividing resources, but there’s a twist. At every time-step, the revenue shares received at time t are agent resources at time $t + 1$, and the game is repeated. The key issue here is that the way resources are shared has a dramatic effect on long-term social welfare, so in order to maximize individual long-term revenue one must consider the welfare of others, a behavior not captured by other models of cooperation and bargaining. Our work focuses on homogeneous production functions. We identify conditions that ensure that the socially optimal outcome is an ϵ -Nash equilibrium. We apply our results to some families of utility functions, and discuss their strategic implications.

1 Introduction

Consider the following classic problem: a group of agents wishes to divide some jointly earned revenue among themselves in some reasonable manner. The economic literature offers several revenue division methods satisfying some notion of fairness: for example, if the value of agent subsets is given, one can use concepts from coalitional games. If we only know agent utilities, we may use bargaining solutions. However, it is not clear that agents will immediately agree to employing a fair solution concept: if there is no external mediator enforcing the solution, it is reasonable to assume that selfish agents will simply push to obtain as large a share as possible, foregoing any fairness considerations.

The problem here is that greedy agents see their interaction with others as a “one shot” affair; demanding more now does not affect their future welfare (this idea is captured by the famous “tragedy of the commons” problem). However, the greediness assumption seems unreasonable. In many settings, such as budget division in companies or governments, the way revenue is shared at the current round affects welfare in future rounds. For example, the head of a government office desires a high share of the national budget, but, even

if one assumes that she only cares about her ministry’s success, the success of one ministry depends on the success of others: the ministry of education requires roads, electricity, and health services, all provided by other departments. This means that in iterative settings, non-myopic agents — i.e., *selfish agents* who care about their *long-term welfare*, rather than only *immediate* rewards — must balance their individual welfare against the needs of society as a whole.

This type of tension between individual desires and global utility is common in settings where budgets need to be iteratively shared. One can think of departments in a company that need to share a budget or manpower, agents that need to share the costs of a public project among themselves, or researchers competing for funding with their peers.

Agents are faced with two conflicting agendas. On the one hand, each agent wants to maximize its own revenue; on the other hand, taking up a large share may hurt future profits and result in lower future revenue: even the most incompetent minister would see that investing 100% of the country’s budget in a single ministry is detrimental to future growth.

To model this setting, we introduce *strategic negotiation games*: agents work together for several rounds, generating revenue according to some production function, and use the revenue that they jointly generate in order to maintain their operations (e.g. pay salaries or buy equipment). At each round, each agent proposes a way of dividing the revenue generated at that round, and an aggregate of the proposals is used to divide the revenue.

1.1 Our Contributions

The main message of this work is that non-myopic agents can arrive at socially optimal, equilibrium outcomes without the need for external intervention. In other words, agents may disagree on how revenue should be divided; but, granted that they see the long-term merit of others’ contributions, they will strive to keep others productive. To conclude, non-myopic negotiators can actually agree on an outcome, and that outcome will be socially optimal.

We formally prove this in a class of games called *strategic negotiation games*. In these settings, agent strategies are revenue sharing proposals, and their utility functions are their total share of the profits for T rounds. However, our results actually hold for other models of negotiation and collaboration (see Section 3.1).

We show that when the production function v is homogeneous of degree $k \geq 1$, agents start displaying more collaborative behavior as T grows; that is, the socially optimal revenue division is an eventual equilibrium. However, when v is homogeneous of degree $k < 1$, this guarantee does not hold: agent strategies do not result in an outcome that is necessarily socially or individually optimal.

Section 6 discusses the implications of our results to functions that model real-world economic interactions, such as CES, Cobb-Douglas and Leontief production functions, as well as to network flow games.

1.2 Related Work

Equilibria in strategic negotiation games are somewhat similar to farsighted equilibria discussed by Kilgour [1984] and Li [1992]. There, agents engage in a repeated game, and make strategic outcomes based on how they anticipate others will behave in future rounds. Our work differs in a crucial manner: since agent actions affect future revenue, agents are not playing the same game repeatedly, but rather a different game in each round.

Bargaining and negotiation outcomes are well understood (see the overview in [Chevalyre *et al.*, 2006] or [Weiss, 2013, Chapter 4], as well as the works by Clippel *et al.* [2008], Kalai and Smorodinsky [1975], and Nash [1950]). Some works (e.g. [Endriss *et al.*, 2006]) consider the effects of negotiation on social welfare, as well as iterated team formation [Dutta *et al.*, 2005; Konishi and Ray, 2003; Shehory and Kraus, 1996]. Cooperative solution concepts such as the Shapley value or the core [Peleg and Sudhölter, 2007] propose reasonable payoff divisions that consider the value of subsets of agents. Models of dynamics and bargaining in cooperative games do exist (see the seminal work by Aumann and Maschler [1964], Davis and Maschler [1963], and Harsanyi [1963], or [Lehrer and Scarsini, 2012; Elkind *et al.*, 2013] for a literature review), but these models (as do models of bargaining and negotiation) do not consider the effects of revenue distribution on future welfare.

Our work is also similar to the problem of *portfolio selection* [Cover, 1991]. One can think of our setting as one where the individual stocks have incentives and would like more money invested in them.

2 Preliminaries

Throughout the paper, boldface letters refer to vectors; the set \mathbb{N} is the set of all positive integers (excluding 0), and the set \mathbb{R}_+ is the set of all non-negative real numbers (including 0). A group of agents $N = \{1, \dots, n\}$ has an *initial resource vector* $\mathbf{w}(0) = (w_1(0), \dots, w_n(0)) \in \mathbb{R}_+^n$, there is some *production function* $v : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$, which, for every $\mathbf{x} \in \mathbb{R}_+^n$, determines the production level $v(\mathbf{x}) \in \mathbb{R}_+$. We assume that both $\mathbf{w}(0)$ and v are known to all agents.

Time is divided into discrete epochs indexed by t . At time $t = 1$, the total payoff (production value) is $v(\mathbf{w}(0))$, which we denote as V_1 ; agents select an allocation vector \mathbf{x}_1 from the $n - 1$ dimensional simplex ($\mathbf{x}_1 \in \Delta_n$, where $\Delta_n = \{\mathbf{x} \in \mathbb{R}_+^n \mid \sum_{i=1}^n x_i = 1\}$), and each agent $i \in N$ receives an individual share of $w_i(1) = x_{i,1}V_1$.

At time 2, the aggregate production is $V_2 = v(\mathbf{w}(1))$, agents select $\mathbf{x}_2 \in \Delta_n$, and the process repeats.

A *contract* χ is a sequence $(\mathbf{x}_t)_{t=1}^\infty$ of vectors in Δ_n . At every time step t , agent i receives a share of $w_i(t) = x_{i,t}V_t$, where $V_t = v(\mathbf{w}(t-1))$. The revenue that a contract χ generates at time t is denoted $V_t(\chi)$, and is computed recursively as described above. Observe that the value $V_t(\chi)$ depends only on the first $t - 1$ payoff divisions; we sometimes write $V_t(\mathbf{x}_1, \dots, \mathbf{x}_{t-1})$ when this property needs to be emphasized.

We say that a contract is *stationary* if for all t we have $\mathbf{x}_t = \mathbf{x}$ for some fixed $\mathbf{x} \in \Delta_n$. A stationary contract is then identified with \mathbf{x} , rather than the constant sequence $(\mathbf{x})_{t=1}^\infty$.

Given a contract $\chi = (\mathbf{x}_t)_{t=1}^\infty$, the *total welfare at time T* is simply $\sum_{t=1}^T V_t(\chi)$, and is denoted $\text{sw}_T(\chi)$. We say that a contract χ^* is *pointwise optimal at time T* if χ^* maximizes $\text{sw}_T(\chi)$ over the space of all possible contracts. We say that χ^* is *universally optimal* if χ^* is optimal at time T for all T .

We assume that agents are not interested in the global social welfare provided by a contract, but rather in their own revenue. Given an agent $i \in N$, the benefit that agent i receives from the contract $\chi = (\mathbf{x}_t)_{t=1}^\infty$ at time t is given by $x_{i,t}V_t(\chi)$; we also write $U_{i,T}(\chi) = \sum_{t=1}^T x_{i,t}V_t(\chi)$, to be agent i 's *utility* (or *individual revenue*) at time T . We mention that other utility functions can be considered, with similar results (see Section 3.1) We say that χ^* is *individually optimal for agent i at time T* if χ^* maximizes $U_{i,T}$ over the space of all possible contracts. Observe that if v is continuous, then for all T there exists some contract χ^* that is individually optimal for i at time T .

3 Strategic Negotiation Games

Strategic negotiation games naturally model budget division and negotiation in iterated settings. Each agent starts with an initial endowment of $w_i(0)$. Just as in Section 2, at time t , $V_t = v(\mathbf{w}(t-1))$ is the amount to be divided among the agents. The only difference is in the way that $\mathbf{w}(t)$ is computed. First, each agent $i \in N$ proposes a sequence of T revenue divisions $\mathbf{y}_1(i), \dots, \mathbf{y}_T(i) \in \Delta_n$. Agents aggregate their proposals by taking an average: $\mathbf{x}_t = \sum_{i=1}^n \frac{1}{n} \mathbf{y}_t(i)$.¹ At round 1, agents divide $V_1 = v(\mathbf{w}(0))$; if they arrive at the aggregate contract $\mathbf{x}_1 \in \Delta_n$, the resources for the next round are $\mathbf{w}(1) = (x_{1,1}V_1, \dots, x_{n,1}V_1)$, and $V_2 = v(\mathbf{w}(1))$; similarly, the game is repeated on round $t + 1$ with $w_i(t) = x_{i,t}V_t$, and $V_{t+1} = v(\mathbf{w}(t))$. Agent i 's utility is given by $U_{i,T} = \sum_{t=1}^T w_i(t)$. The value T is referred to as the *horizon* of the game; agents are said to *see T steps ahead*.

In this setting, agents can only gain by proposing that they receive 100% of the revenue at time T (in particular, if $T = 1$ demanding 100% of the revenue is a dominant strategy); however, in Section 5, we show that for a large enough value of T , the contracts that agents propose at earlier rounds are close to optimal (i.e., nearly maximize social welfare).

¹One can assume more complex aggregation functions rather than a simple average. Our results hold whenever the aggregation function f is continuous, and satisfies a notion of unanimity (i.e., if all agents propose \mathbf{x} , the aggregate outputs \mathbf{x}).

In the static variant of this game, agents propose a single contract $\mathbf{y}(i)$ to be used in all T rounds. We refer to this case as a strategic negotiation game with *static contracts*.

Before we proceed with our formal results, let us demonstrate what happens when agents' behavior is greedy.

Example 3.1. Suppose that we have two agents, and a production function $v(x, y) = \alpha x^2 y^3$, where $\alpha = \left(\frac{5}{3}\right)^3 \cdot \left(\frac{5}{2}\right)^2$. We assume that agents' starting resources are $w_1(0) = \frac{2}{5}$, $w_2(0) = \frac{3}{5}$. If agents' utilities are just their revenue at time t , then their proposals will be $(1, 0)$ and $(0, 1)$, respectively, arriving at the contract $(\frac{1}{2}, \frac{1}{2})$. If agents choose to divide revenue equally, the resulting payoffs decay rapidly (at $t = 3$, agent 1's revenue is $\frac{1}{2} \cdot \frac{\alpha^{31}}{2^{155}}$); however, the socially optimal contract $(\frac{2}{5}, \frac{3}{5})$ offers them a constant revenue at each round. Thus, adopting the optimal contract results in strictly higher revenue in the long run.

Example 3.1 demonstrates the issue with the greedy behavior: if agents are non-myopic, their proposals result in higher revenue both for them and for the group as a whole. In Section 5, we formally explain this intuition.

Before we proceed, let us recall that an ε -Nash equilibrium of a game is a strategy profile where no agent can gain more than ε by changing its action. Thus, in the strategic negotiation setting, agent i is not interested in changing his proposed share if doing so will not increase his utility by more than ε . We mention that similar results hold for other solution concepts (see Section 3.1). Since we want to capture the effects of the horizon T on the equilibria, it will be useful to think of some vectors as *eventual consensus points*.

Definition 3.2 (Eventual Consensus Point). Given N , $\mathbf{w}(0) \in \mathbb{R}_+^n$ and $v : \mathbb{R}_+ \rightarrow \mathbb{R}$, a point $\mathbf{x}^* \in \Delta_n$ is an eventual consensus point of $\langle N, \mathbf{w}(0), v \rangle$ at time t if for every $\varepsilon > 0$ there is some T_0 such that for all $T > T_0$, if all agents see T steps ahead, there is an ε -Nash equilibrium of the resulting negotiation game where all agents propose \mathbf{x}^* at time t .

In other words, \mathbf{x}^* is a consensus point if sufficiently farsighted agents will jointly agree to propose it at time t . When we are limited to static contracts, agents must essentially agree on a single revenue division to be used in the next T rounds. The following definition captures an eventual agreement on the contract to propose.

Definition 3.3 (Eventual Static Equilibrium). A point $\mathbf{x}^* \in \Delta_n$ is an eventual static equilibrium of $\langle N, \mathbf{w}(0), v \rangle$ if for every $\varepsilon > 0$ there is some T_0 such that for all $T > T_0$, if all agents see T steps ahead, then all agents proposing \mathbf{x}^* is a symmetric ε -Nash equilibrium of the resulting strategic negotiation game under static contracts.

3.1 Model Variants

Our results still hold if $U_{i,T}$ is only the payoff at time T , rather than the total payoff up to time T .

Strategic negotiation is only one example of a strategic interaction where non-myopic behavior leads to desirable outcomes. Similar results hold in other strategic settings, and

²We have chosen v to be a Cobb-Douglas production function; this is not a coincidence, as we will show in Section 5

with other solution concepts. For example, one can employ the ε -Strong Nash equilibrium.

Our model does not use discounted returns: future revenue is seen as equally important as current revenue; however, we can easily incorporate discount factors, with similar results. In general, agent welfare is far more affected by the structure of the production function than by discounted returns.

To conclude, non-myopic behavior is a widely applicable paradigm, that leads agents to behave in a desirable manner, without the need for external intervention.

4 Optimal Contracts: First Observations

Finding a socially optimal contract at time T amounts to solving the following problem:

$$\begin{aligned} \max \quad & \sum_{t=1}^T v(\mathbf{w}_t) \quad \text{over} \quad (\mathbf{w}_t)_{t=1}^T \\ \text{s.t.} \quad & \sum_{i=1}^n w_{i,1} = V_1 \\ & \sum_{i=1}^n w_{i,t} = v(\mathbf{w}_{t-1}) \quad \forall t, 2 \leq t \leq T \end{aligned} \quad (1)$$

We say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *monotone increasing* if for every $\mathbf{w}, \mathbf{w}' \in \mathbb{R}^n$ such that $\mathbf{w}' \leq \mathbf{w}$ (i.e. $w'_i \leq w_i$ for all i), $f(\mathbf{w}') \leq f(\mathbf{w})$. We first show that if v is monotone increasing, one can find an optimal solution to Equation (1) using a greedy procedure. Given a non-negative constant C , let us write $\Delta_n(C) = \{\mathbf{w} \in \mathbb{R}_+^n \mid \sum_{i=1}^n w_i = C\}$; in particular, Δ_n equals $\Delta_n(1)$.

Proposition 4.1. *Let $\text{opt}(V)$ be the value of Equation (1) when $V_1 = V$; if v is monotone increasing, then $\text{opt}(V)$ is monotone increasing in V .*

As an immediate corollary of Proposition 4.1 (proof omitted), we have that in order to find an optimal contract at time T , one must find a point $\mathbf{w}_1 \in \text{argmax}_{\mathbf{w} \in \Delta_n(V_1)} v(\mathbf{w})$; then, a point $\mathbf{w}_2 \in \text{argmax}_{\mathbf{w} \in \Delta_n(V_2)} v(\mathbf{w})$ and so on. In other words, greedily maximizing the revenue at every time step will result in a universally optimal contract. From a computational point of view, finding universally optimal contracts is relatively straightforward when v is monotone, assuming that maximizing v over $\Delta_n(V)$ can be done in polynomial time for any V .

When v is homogeneous, finding an optimal contract is even easier. A function $f : \mathbb{R}_+^n \rightarrow \mathbb{R}_+$ is *homogeneous of degree k* , or *k -homogeneous*, if for all $\alpha \geq 0$ and all $\mathbf{w} \in \mathbb{R}_+^n$, $v(\alpha \mathbf{w}) = \alpha^k v(\mathbf{w})$. We now make an important observation (proof omitted): if v is homogeneous of degree k , then there exists a stationary socially optimal contract.

Proposition 4.2. *If v is monotone and homogeneous of degree k , then the set of universally optimal stationary contracts is the set of global maxima of v over Δ_n .*

Another useful property of homogeneity is as follows: if v is homogeneous, we get a closed formula for utility under a contract $\chi = (\mathbf{x}_t)_{t=1}^\infty$.

$$V_T(\chi) = V_1^{k^{T-1}} \prod_{t=1}^{T-1} v(\mathbf{x}_t)^{k^{T-1-t}}, \quad (2)$$

and an even simpler formula for stationary contracts:

$$V_T(\mathbf{x}) = V_1^{k^{T-1}} v(\mathbf{x})^{\sum_{t=0}^{T-2} k^t}. \quad (3)$$

Applying Equation (2), the individual utility of agent i at time T is thus $\sum_{t=1}^T x_{i,t} V_1^{k^{t-1}} \prod_{h=1}^{t-1} v(\mathbf{x}_h)^{k^{t-1-h}}$, and for a stationary contract, it is $x_i \left(\sum_{t=1}^T V_1^{k^{t-1}} v(\mathbf{x})^{\sum_{h=0}^{t-2} k^h} \right)$.

5 Optimal Contracts are Eventual Equilibria and Consensus Points

Using the closed formulas for individual agent utility, we now present the main result of this paper: when agents are sufficiently far-sighted, socially optimal contracts are eventual equilibria of strategic negotiation games.

In order to have a robust model of cooperation and individual incentives, we wish to capture some notion of complementarity among agents. If agents are actually better off without allocating resources to some agents, then there is little reason to collaborate with them. This notion is captured via the idea of *mutual dependency*.

Definition 5.1. A function $v : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies mutual dependency if for all $\mathbf{x} \in \Delta_n$, if $x_i = 0$ for some $i \in N$, then $v(\mathbf{x}) = 0$.

Mutual dependency is an important aspect of negotiating optimal contracts; if mutual dependency does not hold, then there are some agents who need not contribute to the group effort, and are somewhat expendable. In other words, suppose that \mathbf{x}^* is an optimal stationary contract for v and that $x_i^* = 0$; then the set $N \setminus \{i\}$ can argue that agent i should receive no share of the profits, as there exist optimal contracts that do not require any of his resources (in terms of a company, i would be a department that is completely redundant).

Remark 5.2. Mutual dependency is a sufficient — not necessary — condition for our results to hold. Many weaker forms of complementarities will suffice: we really require that all maxima (for agents as a whole and as individuals) occur in the interior of Δ_n . In fact, in Section 6, we show that our results apply to CES production functions, for which the mutual dependency does not hold.

Model assumptions to guarantee interior solutions have been used in other settings (e.g. [Kelly *et al.*, 1998; Kelly and Voice, 2005]); weaker versions of our results do hold in the case where maxima occur on the boundary of Δ_n . Briefly, if there are global maxima on the boundary of Δ_n , then in the limit, agents who receive a positive share of the profits will want to have these maxima as the negotiation outcome. That said, it is unclear what strategies can achieve this outcome.

In what follows we assume that the function v satisfies mutual dependency, and that there are some points in Δ_n for which v assumes strictly positive values.

We begin with a basic question regarding properties of optimal versus individually optimal stationary contracts. Given an optimal (or pointwise optimal) contract $\mathbf{x}^* \in \Delta_n$, is it possible that \mathbf{x}^* is also individually optimal for some $i \in N$? As Lemma 5.3 shows (proof omitted), if v is continuously differentiable, this is not possible.

Lemma 5.3. *Suppose that v is continuously differentiable, and define $F : \mathbb{R}^n \rightarrow \mathbb{R}$ with $F \in \{V_t, \text{sw}_t\}$ (where t is a fixed time period); Given some \mathbf{x}^* that is a global maximum*

of F in the interior of Δ_n , then \mathbf{x}^ does not maximize individual utility for any $i \in N$ at time t .*

Lemma 5.3 implies that there always exist stationary contracts that are better for individuals than optimal stationary contracts. In particular, an optimal contract is never a symmetric equilibrium of the strategic negotiation game. The proof of Lemma 5.3 is simple, but we wish to stress its importance. For an individual agent, this lemma presents a rather unfortunate state of affairs: any contract that is socially optimal is *necessarily* individually sub-optimal. This means that social welfare and individual gains are always at odds. Can we, under certain conditions, mitigate this effect? In what follows, we show that socially optimal contracts can be eventual consensus points, and eventual equilibria of strategic negotiation games.

Before we proceed, let us recall the following property of homogeneous functions.

Proposition 5.4. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a k -homogeneous, differentiable function, then the point \mathbf{x}^* is a critical point of f in the interior Δ_n if and only if $\frac{\partial f}{\partial x_i}(\mathbf{x}^*) = \frac{\partial f}{\partial x_j}(\mathbf{x}^*)$ for all $i, j \in N$; in that case, $\frac{\partial f}{\partial x_i}(\mathbf{x}^*) = kf(\mathbf{x}^*)$ for all $i \in N$.*

Using the formulas for individual utility derived from Equations (2) and (3), we obtain the following result.

Lemma 5.5. *Let $\chi_T^* = (\mathbf{x}_t^*)_{t=1}^\infty$ be an individually optimal contract for agent i at time T , then for any $q \leq T$ and any $j, j' \in N$ such that $j, j' \neq i$, $\frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*) = \frac{\partial v}{\partial x_{j',q}}(\mathbf{x}_q^*)$; moreover,*

$$\frac{\partial v}{\partial x_{i,q}}(\mathbf{x}_q^*) = \frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*) - \frac{v(\mathbf{x}_q^*)k^q V_q(\chi_T^*)}{\sum_{t=q}^T k^t x_{i,t} V_t(\chi_T^*)}$$

Proof. First, if χ_T^* is an individually optimal contract at time T , then in particular, $(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*)$ is a critical point of $U_{i,T}(\chi_T^*)$ subject to the constraint that $\mathbf{x}_t^* \in \Delta_n$ for all $1 \leq t \leq T$. Translating this to Lagrange multipliers, we have a function

$$L(\mathbf{x}_1, \dots, \mathbf{x}_T, \lambda_1, \dots, \lambda_T),$$

that equals $U_{i,T}(\mathbf{x}_1, \dots, \mathbf{x}_T) + \sum_{t=1}^T \lambda_t (1 - \sum_{j=1}^n x_{j,t})$. Since $(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*)$ are a critical point of $U_{i,T}$ subject to their constraints, we get that for any $1 \leq q \leq T$ and any $j \neq i$,

$$\frac{\partial L}{\partial x_{j,q}}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*, \lambda_1, \dots, \lambda_T) = \frac{\partial U_{i,T}}{\partial x_{j,q}} - \lambda_q = 0.$$

Now,

$$\frac{\partial U_{i,T}}{\partial x_{j,q}}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*) = \sum_{t=1}^T x_{i,t}^* \frac{\partial V_t(\mathbf{x}_1^*, \dots, \mathbf{x}_t^*)}{\partial x_{j,q}} \quad (4)$$

which, by the closed-form formula for V_t , equals

$$\sum_{t=q}^T x_{i,t}^* V_1^{k^t} \left(\prod_{\substack{h=1 \\ h \neq q}}^t v(\mathbf{x}_h^*)^{k^{t-h}} \right) k^{t-q} v(\mathbf{x}_q^*)^{k^{t-q-1}} \frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*),$$

which we simplify to $\frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*) \frac{1}{k^q v(\mathbf{x}_q^*)} \sum_{t=q}^T x_{i,t}^* k^t V_t(\chi_T^*)$. Note that $v(\mathbf{x}_q^*) > 0$ since χ_T^* is individually optimal for i

at time T , and in particular, i always chooses a point where positive profit is made, due to mutual dependency. Since $\frac{\partial U_{i,T}}{\partial x_{j,q}}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*) = \frac{\partial U_{i,T}}{\partial x_{j',q}}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*)$ for all $j, j' \neq i$, we get that $\frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*) = \frac{\partial v}{\partial x_{j',q}}(\mathbf{x}_q^*)$ for all $j, j' \neq i$ and for all $q \leq T$. Similarly, we take the derivative of $U_{i,T}$ with respect to $x_{i,q}$ to obtain

$$\frac{\partial U_{i,T}}{\partial x_{i,q}}(\mathbf{x}_1^*, \dots, \mathbf{x}_T^*) = V_q(\chi_T^*) + \sum_{t=1}^T x_{i,t}^* \frac{\partial V_t(\mathbf{x}_1^*, \dots, \mathbf{x}_t^*)}{\partial x_{i,q}}$$

which equals

$$V_q(\chi_T^*) + \frac{\partial v}{\partial x_{i,q}}(\mathbf{x}_q^*) \frac{1}{k^q v(\mathbf{x}_q^*)} \sum_{t=q}^T x_{i,t}^* k^t V_t(\chi_T^*)$$

which implies that for all $j \neq i$,

$$\frac{\partial v}{\partial x_{i,q}}(\mathbf{x}_q^*) = \frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q^*) - \frac{V_q(\chi_T^*) k^q v(\mathbf{x}_q^*)}{\sum_{t=q}^T x_{i,t}^* k^t V_t(\chi_T^*)}.$$

□

Lemma 5.5 implies that if $\chi_T^* = (\mathbf{x}_t^*)_{t=1}^\infty$ is an individually optimal contract for $i \in N$ at time T , then every point \mathbf{x}_q^* is “nearly” a critical point of v : the partial derivatives of v at \mathbf{x}_q^* are all equal, except that of agent i . However, when fixing a time q and taking the horizon T to infinity, one can ensure that \mathbf{x}_q^* approaches a critical point of v over Δ_n .

First, let us observe a simple property of individually optimal contracts.

Lemma 5.6. *Suppose that v is k homogeneous. Let $\chi_T^* = (\mathbf{x}_t^*)_{t=1}^\infty$ be an individually optimal contract for agent i at time T , and let \mathbf{x}^* be an optimal stationary contract, then $x_{i,t}^* \geq x_i^*$ for all t .*

Proof. Assume otherwise, then at time t agent i receives a share of $x_{i,t}^* v(\mathbf{x}_t^* V_{t-1}(\chi_T^*))$. Under \mathbf{x}^* , agent i could receive at time t $x_i^* v(\mathbf{x}_t^* V_{t-1}(\chi_T^*))$, which is strictly greater than what he receives at time t under χ_T^* . Moreover, since \mathbf{x}^* maximizes v over Δ_n , it also maximizes v over $\Delta_n(V_{t-1}(\chi_T^*))$; i.e., choosing \mathbf{x}^* will strictly improve welfare for i at time t , and weakly improve his welfare at every subsequent round. □

Since Lemma 5.6 holds for any individually optimal contract, we get as a corollary that there is some constant $c > 0$ such that $x_{i,t}^* > c$ for all individually optimal contracts and for all $t \in \mathbb{N}$. Given a contract $\chi = (\mathbf{x}_t)_{t=1}^\infty$, agent i receives a weakly increasing share at every round if for all t ,

$$x_{i,t} V_t(\chi) \leq x_{i,t+1} V_{t+1}(\chi).$$

Theorem 5.7. *Suppose that we are given a sequence of contracts $(\chi_T^*)_{T=1}^\infty$ such that χ_T^* is individually optimal for agent i at time T for all T . Moreover, suppose that i receives a weakly increasing share under χ_T^* for all T .*

If v is continuously differentiable and homogeneous of degree $k \geq 1$, then for any $q \in \mathbb{N}$, if $(\mathbf{x}_q(T))_{T=1}^\infty$ converges, then $\lim_{T \rightarrow \infty} \mathbf{x}_q(T)$ is a critical point of v over Δ_n . However, if $k < 1$, then the limit is never a critical point.

Proof. The case $k < 1$ is omitted due to space constraints.

We show that when $k \geq 1$, we do have convergence to a critical point. According to Lemma 5.5,

$$\frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q(T)) = \frac{\partial v}{\partial x_{j',q}}(\mathbf{x}_q(T))$$

for all $j, j' \neq i$. Now, it remains to show that

$$\lim_{T \rightarrow \infty} \frac{\partial v}{\partial x_{i,q}}(\mathbf{x}_q(T)) - \frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q(T)) = 0.$$

Under the above assumptions, this indeed holds. All we need to show is that

$$\lim_{T \rightarrow \infty} \frac{V_q(\chi_T^*) k^q v(\mathbf{x}_q(T))}{\sum_{t=q}^T x_{i,t}(T) k^t V_t(\chi_T^*)} = 0.$$

Letting $\mathbf{x}^* \in \operatorname{argmax}_{\mathbf{x} \in \Delta_n} v(\mathbf{x})$ be some optimal stationary contract, we recall that according to Lemma 5.6, $x_{i,q}(T) \geq x_i^*$, and that $x_i^* > 0$ according to the mutual dependency property. Therefore

$$\begin{aligned} \frac{V_q(\chi_T^*) k^q v(\mathbf{x}_q(T))}{\sum_{t=q}^T x_{i,t}(T) k^t V_t(\chi_T^*)} &\leq \frac{V_q(\chi_T^*) k^q v(\mathbf{x}_q(T))}{\sum_{t=q}^T x_{i,q}(T) k^t V_q(\chi_T^*)} \\ &= \left(\frac{k^q v(\mathbf{x}_q(T))}{x_{i,q}(T)} \right) \frac{1}{\sum_{t=q}^T k^t} \\ &\leq \frac{v(\mathbf{x}^*)}{x_i^*} \frac{1}{\sum_{t=0}^{T-q} k^t} \end{aligned}$$

Since v is homogeneous of degree $k \geq 1$, we have that $\sum_{t=0}^{T-q} k^t$ goes to infinity as T grows, thus

$$\lim_{T \rightarrow \infty} \frac{\partial v}{\partial x_{i,q}}(\mathbf{x}_q(T)) - \frac{\partial v}{\partial x_{j,q}}(\mathbf{x}_q(T)) = 0.$$

According to Proposition 5.4, this implies that the limit of the sequence $(\mathbf{x}_q(T))_{T=1}^\infty$ is a critical point of v over Δ_n . □

If v has a *unique* critical point in the interior of Δ_n (e.g., if v is a strictly concave function), then we can obtain a stronger claim.

Corollary 5.8. *Let $(\chi_T^* = (\mathbf{x}_t(T))_{t=1}^\infty)_{T=1}^\infty$ be a sequence satisfying the conditions stated in Theorem 5.7; suppose that v is continuously differentiable, concave and homogeneous of degree $k \geq 1$; then for any q , $\lim_{T \rightarrow \infty} \mathbf{x}_q(T)$ exists and is an optimal contract for v .*

Corollary 5.8 implies the following claim

Theorem 5.9. *If v is continuously differentiable, homogeneous of degree $k \geq 1$, and has a unique maximum \mathbf{x}^* over Δ_n , then for every fixed $q \in \mathbb{N}$, \mathbf{x}^* is an eventual consensus point of $\langle N, \mathbf{w}(0), v \rangle$ for all $t \leq q$.*

Simply put, Theorem 5.9 states that as the horizon increases, agents can agree on sharing according to the optimal contract on early rounds; however, there is no such guarantee for later rounds.

In strategic negotiation games with static contracts, we can show an even stronger result: if we assume that the socially optimal contract \mathbf{x}^* can secure non-vanishing revenue to every one of the agents — i.e. that $\lim_{T \rightarrow \infty} U_{i,T}(\mathbf{x}^*) = \infty$ for all $i \in N$ — then \mathbf{x}^* is an eventual consensus point for all t .

Theorem 5.10. *Suppose that v is homogeneous of degree $k \geq 1$ and has a unique maximum \mathbf{x}^* over Δ_n ; if agents have non-vanishing revenue under \mathbf{x}^* , then \mathbf{x}^* is an eventual static equilibrium of $\langle N, \mathbf{w}(0), v \rangle$.*

Theorems 5.9 and 5.10 do not hold if we drop the unique maximum assumption. Let $\mathbf{x}^*, \mathbf{y}^* \in \Delta_n$ be two different maxima of v in Δ_n . There exist $i, j \in \Delta_n$ such that $x_i^* > x_j^*$ and $y_i^* < y_j^*$. Under \mathbf{x}^* , agent j will do much better, and similarly, under \mathbf{y}^* , agent i will do much better. In this case, our results imply that for every $i \in N$, there is a socially optimal contract that is “nearly” individually optimal for i .

Theorems 5.9 and 5.10 rely on v being continuously differentiable; we can, prove similar claims for functions that are uniform limits of continuously differentiable functions. Such functions are not mere pathologies, but arise in some of the applications we describe later, in Section 6. This is stated in Theorem 5.11 (proof omitted).

Theorem 5.11. *If v is the uniform limit of continuously differentiable functions that are homogeneous of degree $k \geq 1$, and v has a unique maximum \mathbf{x}^* over Δ_n , then for any fixed $q \in \mathbb{N}$ \mathbf{x}^* is an eventual consensus point for all $t \leq q$. Moreover, \mathbf{x}^* is an eventual static equilibrium of $\langle N, \mathbf{w}(0), v \rangle$ under static contracts.*

6 Applications

In this section we analyze certain classes of functions to which our results apply. In first set of examples, we explore common production functions used in the economic literature. The first three functions we study have inherent agent parameters, a_1, \dots, a_n . We also write $a = \sum_{i=1}^n a_i$.

We first look at *CES production functions* (Constant Elasticity of Substitution); for two agents, constant elasticity of substitution captures the idea that an increase in the good of one agent would cause a constant reduction in the amount of good required by the other, in order to produce the same amount of revenue. Let $v_r(\mathbf{x})$ be $c \cdot \left(\frac{1}{n} \sum_{i=1}^n \left(\frac{x_i}{a_i} \right)^r \right)^{\frac{1}{r}}$, where c and a_1, \dots, a_n are positive constants, and $r \neq 0$. When $r < 1$, v_r is concave and homogeneous of degree 1. We observe that in general, CES production functions do not satisfy the mutual dependency property. For example $v(x, y) = (\sqrt{x} + \sqrt{y})^2$ is a CES production which satisfies $v(1, 0) = v(0, 1) = 1$. However, when $r < 1$, CES production functions are homogeneous of degree 1, and have a unique maximum over Δ_n , and individual players’ maxima occur in interior points (when the horizon T is more than 1).

A *Cobb-Douglas production function* is a function of the form $v_c(\mathbf{x}) = c \prod_{i=1}^n x_i^{a_i}$. Cobb-Douglas production functions also capture some notion of substitutability among agents; the variable exponents strongly determine the payment to agent i under an optimal contract. Indeed, we observe that a Cobb-Douglas production function is a -homogeneous: $v_c(\lambda \mathbf{x}) = c \prod_{i=1}^n (\lambda x_i)^{a_i} = \lambda^a v_c(\mathbf{x})$. It is well-known in the economic literature that the maximum of v_c over Δ_n is unique, and equals $(\frac{a_1}{a}, \dots, \frac{a_n}{a})$, with the shares directly proportional to the weights a_i . Moreover, the Cobb-Douglas production function is concave.

Next, consider the function $v_\ell(\mathbf{x}) = c \min_{i \in N} \left\{ \frac{x_i}{a_i} \right\}$, known as a *Leontief production function*, where c and $(a_i)_{i \in N}$ are all strictly positive constants. If agent production is determined by Leontief functions, their goods are entirely unsubstitutable. For example, in order to produce lemonade (priced at \$2/cup), one needs 4 cups of lemon juice (priced at \$0.1/cup), 1 cup of sugar (priced at \$0.02/cup) and half a cup of mint (priced at \$0.5/cup). A Leontief production function capturing the revenue that agents can produce given that they invest x, y, z dollars in lemon juice, sugar and mint, respectively, is

$$v_\ell(x, y, z) = \$1 \times \min \left\{ \frac{x}{4 \times \$0.1}, \frac{y}{1 \times \$0.02}, \frac{z}{0.5 \times \$0.5} \right\}.$$

The function v_ℓ is 1-homogeneous. However, since v_ℓ is not differentiable, Theorem 5.7 does not apply here. However, $\lim_{r \rightarrow -\infty} v_r(\mathbf{x}) = v_\ell(\mathbf{x})$; moreover, v_ℓ has a unique maximum over Δ_n ($x_i^* = \frac{a_i}{a}$ for all $i \in N$), and is the uniform limit of functions satisfying the conditions of Theorem 5.11 (CES production functions with $r < 1$); thus, Theorem 5.11 applies to Leontief functions.

Finally, suppose we are given a directed, weighted graph $\Gamma = \langle V, E \rangle$, with a source-terminal node pair $s, t \in V$, where the edge set E (with $E = \{e_1, \dots, e_n\}$) and the weight of the edge $e_i \in E$ is a positive integer w_i . Given a vector $\mathbf{x} \in \mathbb{R}^n$, the maximum flow function $v_\Gamma(\mathbf{x})$ is the maximum flow from s to t that can be achieved on Γ when the weight of the i -th edge is given by $x_i w_i$. The edge e_i uses the amount of money it has, x_i , to purchase capacity, which is multiplied by the factor w_i . Given $\mathbf{x} \in \mathbb{R}_+^n$, we write $\Gamma(\mathbf{x})$ to be the graph Γ with capacities $w_i x_i$; thus $\Gamma(\mathbf{1}^n) = \Gamma$, and we indeed assume that $\mathbf{w}(0) = \mathbf{1}^n$. This means that $v_\Gamma(\mathbf{x})$ equals the maximum flow through $\Gamma(\mathbf{x})$. This is a straightforward generalization of network flow games [Peleg and Sudhölter, 2007].

The first observation we make is that v_Γ is homogeneous of degree 1: changing all edge capacities by a factor of λ results in a change of λ to the maximum flow. Interestingly, we can also show that the maximum flow function is the uniform limit of 1-homogeneous, concave, differentiable functions. However, it is only when v_Γ has a unique maximum (i.e., a unique maximum flow in Γ) that our results hold.

Theorem 6.1. *the maximum flow function v_Γ is the uniform limit of 1-homogeneous, concave, differentiable functions.*

Proof. Let us write $\Gamma = \langle V, E \rangle$, where $E = \{e_1, \dots, e_n\}$ is the set of directed edges in Γ . We let w_i be the capacity of the edge e_i . Let $\mathcal{C}(\Gamma)$ be the set of (s, t) cuts of Γ . Given a cut $C \in \mathcal{C}(\Gamma)$, we can write C as a vector \mathbf{w}_C in \mathbb{R}^n , with $w_{i,C} = w_i$ for all $e_i \in C$, and $w_{i,C} = 0$ otherwise. We let M_Γ be an $|\mathcal{C}(\Gamma)| \times n$ matrix whose rows are the vectors \mathbf{w}_C . This means that $v_\Gamma(\mathbf{x})$ can be rewritten as $v_\Gamma(\mathbf{x}) = \min_{C \in \mathcal{C}(\Gamma)} \mathbf{w}_C \cdot \mathbf{x}$. Now, given a vector \mathbf{y} in $\mathbb{R}^{|\mathcal{C}|}$, we let $f_r(\mathbf{y})$ be a CES production function, i.e. $f_r(\mathbf{y}) = c \left(\sum_{j=1}^{|\mathcal{C}|} \frac{1}{|\mathcal{C}|} y_j^r \right)^{\frac{1}{r}}$ (here, $a_1 = \dots = a_{|\mathcal{C}|} = 1$). We write $v_{\Gamma,r} : \mathbb{R}^n \rightarrow \mathbb{R}$ to be $v_{\Gamma,r}(\mathbf{x}) = f_r(M_\Gamma \mathbf{x})$, and observe that $\lim_{r \rightarrow -\infty} v_{\Gamma,r}(\mathbf{x}) = v_\Gamma(\mathbf{x})$ with uniform convergence.

We note that all $v_{\Gamma,r}$ are homogeneous of degree 1, differentiable in the interior of Δ_n , and when $r < 1$, are strictly concave. This concludes the proof. \square

To conclude, if v is a

- (a) CES production function with parameter $r < 1$ or
 - (b) Cobb-Douglas production function with $a \geq 1$ or
 - (c) Leontief function or
 - (d) network flow function (Γ has a unique maximum flow),
- and \mathbf{x}^* is a global maximum of v over Δ_n , then for any fixed $q \in \mathbb{N}$, \mathbf{x}^* is an eventual consensus point of $\langle N, \mathbf{w}(0), v \rangle$ at time $t \leq q$. Moreover, \mathbf{x}^* is an eventual static equilibrium of $\langle N, \mathbf{w}(0), v \rangle$ under static contracts.

7 Conclusions and Future Work

In this paper, we propose a new paradigm: selfish agents that care about their long-term revenue can achieve good outcomes without the help of an outside mediator. This approach can be naturally applied to any strategic setting that can be expressed as a function $v : \mathbb{R}^n \rightarrow \mathbb{R}$. This includes, for example, cooperative games such as matching games, market exchange games, weighted voting games, and weighted graph games.

Not all of these settings induce homogeneous production functions, but we hope that when restricted to classes of games, structural properties of the game can be utilized in order to obtain some interesting results. We expect that in settings where homogeneity does not hold, individuals tend to be less collaborative.

We assume that at every round, agents “reinvest” all of the money towards future production. While this is a realistic assumption in some settings (e.g. government offices or company departments), one could relax this assumption by allowing agents to keep a certain fraction of their revenue at each round. If this share is fixed, our results easily hold. Complicated behavior may arise when agents are allowed to choose how much to invest in the group production and how much to keep for themselves.

Finally, one can naturally extend our model to uncertain environments. A noisy production function would be more realistic; moreover, due to the fact that notions such as mutual dependency may only hold in expectation. In this case, some agents may be paid at a round even if they do not contribute, since they may be useful in future iterations.

Acknowledgments: Part of this work was done while Yair Zick was affiliated with Nanyang Technological University, funded by the A*STAR (SINGA) scholarship.

References

[Aumann and Maschler, 1964] R.J. Aumann and M. Maschler. The bargaining set for cooperative games. In *Advances in Game Theory*, pages 443–476. 1964.

[Chevalleyre *et al.*, 2006] Y. Chevalleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lematre, N. Maudet, J. Padget, S. Phelps, J.A. Rodriguez-Aguilar, and P. Sousa. Issues in

multiagent resource allocation. *Informatica*, 30(1):1–32, 2006.

[Clippel *et al.*, 2008] G. De Clippel, H. Moulin, and N. Tideman. Impartial division of a dollar. *Journal of Economic Theory*, 139(1):176–191, 2008.

[Cover, 1991] T.M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, 1991.

[Davis and Maschler, 1963] M. Davis and M. Maschler. Existence of stable payoff configurations for cooperative games. *Bulletin of the AMS*, 69:106–108, 1963.

[Dutta *et al.*, 2005] B. Dutta, S. Ghosal, and D. Ray. Far-sighted network formation. *Journal of Economic Theory*, 122(2):143–164, 2005.

[Elkind *et al.*, 2013] E. Elkind, D. Pasechnick, and Y. Zick. Dynamic weighted voting games. In *AAMAS’13*, pages 515–522, 2013.

[Endriss *et al.*, 2006] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *JAIR*, 25:315–348, 2006.

[Harsanyi, 1963] J. C. Harsanyi. A simplified bargaining model for the n-person cooperative game. *International Economic Review*, 4(2):194–220, 1963.

[Kalai and Smorodinsky, 1975] E. Kalai and M. Smorodinsky. Other solutions to Nash’s bargaining problem. *Econometrica*, 43(3):513–518, 1975.

[Kelly and Voice, 2005] F. Kelly and T. Voice. Stability of end-to-end algorithms for joint routing and rate control. *Computer Communication Review*, 35(2):5–12, 2005.

[Kelly *et al.*, 1998] F.P. Kelly, A.K. Maulloo, and D.K.H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research society*, pages 237–252, 1998.

[Kilgour, 1984] D.M. Kilgour. Equilibria for far-sighted players. *Theory and Decision*, 16(2):135–157, 1984.

[Konishi and Ray, 2003] H. Konishi and D. Ray. Coalition formation as a dynamic process. *Journal of Economic Theory*, 110(1):1–41, 2003.

[Lehrer and Scarsini, 2012] E. Lehrer and M. Scarsini. On the core of dynamic cooperative games. *CoRR*, abs/1203.2832, 2012.

[Li, 1992] Shuhe Li. Far-sighted strong equilibrium and oligopoly. *Economics Letters*, 40(1):39–44, 1992.

[Nash, 1950] J.F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[Peleg and Sudhölter, 2007] B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*, volume 34. Springer, 2nd edition, 2007.

[Shehory and Kraus, 1996] O. Shehory and S. Kraus. Formation of overlapping coalitions for precedence-ordered task-execution among autonomous agents. In *ICMAS’96*, pages 330–337, 1996.

[Weiss, 2013] G. Weiss, editor. *Multi Agent Systems*. MIT Press, 2013.