A Distributed Mechanism for Public goods Allocation with Dynamic Learning Guarantees

[Short talk]

Abhinav Sinha
University of Michigan
Ann Arbor, Michigan 48109
absi@umich.edu

Achilleas Anastasopoulos
University of Michigan
Ann Arbor, Michigan 48109
anastas@umich.edu

ABSTRACT

In this paper we consider the public goods resource allocation problem (also known as Lindahl allocation) of determining the level of an infinitely divisible public good with P features, that is shared between strategic agents. We present an efficient mechanism, i.e., a mechanism that produces a unique Nash equilibrium (NE), with the corresponding allocation at NE being the social welfare maximizing allocation and taxes at NE being budget-balanced. The main contribution of this paper is that the designed mechanism has two properties, which have not been addressed together in the literature, and aim to make it practically implementable. First, we assume that agents can communicate only through a given network and thus the designed mechanism obeys the agents' informational constraints. This means that each agent's outcome through the mechanism can be determined by only the messages of his/her neighbors. Second, it is guaranteed that agents can learn the NE induced by the mechanism through repeated play when each agent selects a learning strategy from within the "adaptive best-response" dynamics class. This is a class of adaptive learning strategies that includes well-known dynamics such as Cournot best-response, $k-$period best-response and fictitious play, among others. The convergence result is a consequence of the fact that the best-response of the induced game is a contraction mapping. Finally, we present a numerical study of convergence to NE, for two different underlying communication graphs and two different learning dynamics within the ABR class.

CCS CONCEPTS

• **Theory of computation** → **Algorithmic game theory and mechanism design**; **Convergence and learning in games**; **Network games**;

KEYWORDS

Lindahl allocation, Stability of Nash equilibrium, Full implementation, Distributed communication, Fictitious Play, Cournot Learning

This work is supported in part by NSF grant ECCS-1608361.

NetEcon'17, June 27, 2017, Cambridge, MA, USA

ACM ISBN 978-1-4503-5089-1/17/06…$15.00
https://doi.org/10.1145/3106723.3106725

ACM Reference format:
Abhinav Sinha and Achilleas Anastasopoulos. 2017. A Distributed Mechanism for Public goods Allocation with Dynamic Learning Guarantees. In *Proceedings of NetEcon'17, Cambridge, MA, USA, June 27, 2017,* 6 pages.
https://doi.org/10.1145/3106723.3106725

1 INTRODUCTION

The framework of Mechanism design aims to bridge the informational gap between a designer, who wishes to achieve "efficient" allocation and agents, who are strategic and posses private information relevant to determining the efficient allocation. The basic premise of mechanism design has been applied to a variety of applications and as a result, recent works have focused on designing mechanisms that are also practically implementable.

In this vein, there are two very important features of practical mechanisms that haven't been addressed in the literature. (a) The first is the informational constraint between agents - most mechanisms define the contract (e.g., allocation and tax) such that messages from all agents need to be collected centrally in order to determine each agents' outcome. This is akin to assuming a broadcast structure of communication between agents. (b) The second is related to the fact that, when Nash equilibrium (NE) is used as the solution concept, there is little to no guarantee that agents can learn about each others' private information in order to calculate the NE. More specifically, dynamic learning guarantees on convergence to NE induced by the mechanism are typically provided for specific learning dynamics, which can vary from model to model. Addressing the two features together can make a mechanism ready for application to scenarios where agents are distributed - both physically and informationally.

The motivation for feature (a), i.e., designing a distributed mechanism comes from the literature on "distributed optimization", [3, 8, 17, 19], and the motivation for feature (b) comes from the literature on "learning in games", [9, 11, 14, 15, 22].

Regarding problem (a), the literature on distributed optimization aims to address the informational constraints between **non-strategic** agents who possess local private information relevant to the centralized optimization. This model, in-part, is analogous to mechanism design, where in many cases efficient allocation is indeed described by a centralized optimization. Thus a natural question to ask is whether efficient mechanisms can be designed such that they obey the informational constraints of distributed agents. The fact that mechanisms have to deal with strategic agents, means that this is a new and non-trivial question.

It is important to note here that the public goods centralized problem (see (4)) itself is completely oblivious about the informational constraints that the mechanism designer faces. Resolving problem (a) for the public goods problem is not to be confused with the problem of provisioning **local** public goods [1, 21]. In such models, each agents' (intrinsic) utility is assumed to depend only on his/her neighbors allocation thereby naturally aligning the informational constraints with the utility structure. Utilities for the public goods problem do not require any such assumption. Instead, we require that the exchanged messages implied by the mechanism satisfy the network communication constraints.

Regarding problem (b), learning in games is motivated by the fact that NE is, theoretically, a complete information solution concept. Use of NE in models that don't necessarily assume perfect information among agents is typically justified[1] by showing that under certain natural learning strategies, such as evolutionary dynamics [2], or under specifically designed learning strategies, such as regret-minimizing algorithms[2], agents are guaranteed to learn each others' private information and in turn the NE. The larger the class of learning dynamics for which convergence can be guaranteed, the stronger the implied justification.

In this paper, our objective as the designer is to implement the social welfare maximizing allocation for the public goods problem (Lindahl allocation) in the presence of strategic agents. To achieve this we design a mechanism which induces a unique and efficient NE that incorporates both the aforementioned features. Thus, the mechanism determines outcomes for each agent using only his/her neighbors' messages and doesn't require collecting all agents' messages centrally. Also, it is shown that for a sufficiently large class of learning dynamics - *adaptive best-response* dynamics (ABR) - there is guaranteed convergence to NE. The ABR class contains several well-known learning dynamics such as "Cournot best-response" and "fictitious play", [4], among others. Furthermore, the induced game is supermodular and hence there is guaranteed convergence for the *adaptive dynamics* (AD) class, [14], of learning strategies as well.

Related works in the literature that focus on learning in games are as follows. In their seminal work [14], Milgrom and Roberts show that for a *supermodular* game any learning dynamic within the AD class is guaranteed to converge between the two most extreme Nash equilibria. Chen in [6] then presents a mechanism (without any informational constraints between agents) for the Lindahl allocation problem such that the induced game is supermodular and has a unique NE. Following this development, Healy and Mathevet in [10] show that for a *contractive* game, all learning dynamics within the ABR dynamics class converge to the unique NE. These authors also present a mechanism for the Lindahl allocation problem (again without any informational constraints between agents) such that the induced game is contractive. As contraction is a more stringent property than supermodularity, ABR class is broader than the AD class.

[1] Learning of NE can also be justified with the *Evolutive* interpretation of NE, [16, 18], which implies that NE in a single-shot game can be thought of as the stationary point of a dynamic adjustment process.

[2] For instance, for the case of zero-sum games, Daskalakis, Deckelbaum and Kim in [7] propose a new learning algorithm based on regret minimization and show that it converges at a linear rate.

In general, "adaptive" learning strategies are broadly a class of learning strategies where, at each time, agents respond optimally to some combination of empirical distribution arising out of the past observed actions. For example, fictitious play is an adaptive learning strategy. Hofbauer and Sandholm in [11] provide convergence results specifically for fictitious play in the case of zero-sum games, supermodular games and potential games. Monderer and Shapley in [15] provide convergence results for fictitious play in identical interests games, i.e., where best-response is equivalent to that of a game where agents have the same utility functions.

The structure of the remainder of this paper is as follows: Section 2 describes the public goods centralized allocation problem and its optimality conditions. Section 3 defines some mechanism design basics and then presents the mechanism. Section 4 introduces learning-related properties and contains the guaranteed convergence result of any learning dynamic within the ABR class. Finally, Section 4 also contains a numerical study of the convergence pattern for two different learning dynamics and two different underlying communication graphs.

2 THE PUBLIC GOODS PROBLEM

There are N strategic agents, denoted by the set $\mathcal{N} = \{1, \dots, N\}$. A directed communication graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ is given, where the vertexes correspond to the agents and an edge from vertex i to j indicates that agent i can "listen" to agent j. It is assumed that the given graph $\mathcal{G}$ is strongly connected. In this paper, we are interested in the public goods allocation problem, which in Economics literature is also known as Lindahl allocation [12, 13].

There is a single infinitely divisible public good with P features, with the set of features denoted by $\mathcal{P} = \{1, \dots, P\}$. Each agent receives a utility $v_i(x)$ based on the quantity of the public good $x = \left(x^p\right)_{p \in \mathcal{P}} \in \mathbb{R}^P$. Since for each agent, its utility depends on the common allocation x, this is the public goods model. It is assumed that $v_i : \mathbb{R}^P \to \mathbb{R}$ is a continuously double-differentiable, strictly concave function that satisfies, $\forall\, p \in \mathcal{P}$,

$$-\eta < H^{-1}_{pp} + \sum_{l \in \mathcal{P},\, l \neq p} \left| H^{-1}_{pl} \right| < 0, \tag{1a}$$

$$H^{-1}_{pp} < -\frac{1}{\eta}, \tag{1b}$$

for any given $\eta > 1$, where H^{-1} is the inverse of the Hessian $H = \left[(\partial^2 v_i(\cdot)) / (\partial x^k \partial x^l) \right]_{k,l}$. To understand the significance of this assumption consider the case of $P = 1$, then this condition is the same as[3]

$$v_i''(\cdot) \in \left(-\eta, -\frac{1}{\eta} \right). \tag{3}$$

It is already assumed that $v_i(\cdot)$ is strictly concave, the only additional imposition made by this assumption is that the second derivative of $v_i(\cdot)$ is strictly bounded away from 0 and $-\infty$.

The above mentioned properties of the utility function are assumed to be common knowledge between agents and the designer.

[3] More generally if the utility is separable, $v_i(x) = \sum_{p \in \mathcal{P}} v_{i,p}(x^p)$, then the condition in (1) is the same as

$$v''_{i,p}(\cdot) \in \left(-\eta, -\frac{1}{\eta} \right), \quad \forall\, p \in \mathcal{P}. \tag{2}$$

However, the utility function $v_i(\cdot)$ itself is known only to agent i and is not known to other agents or the designer. The designer wishes to allocate the public good such that the sum of utilities is maximized, i.e., to solve the following centralized allocation problem,

$$x^* = \underset{x \in \mathbb{R}^P}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} v_i(x). \tag{4}$$

The allocation x^* is also called the *efficient* allocation and it is assumed to be finite. It is unique due to strictly concave utilities and thus the necessary and sufficient optimality conditions are

$$\frac{\partial v_i(x^*)}{\partial x^p} = \mu_i^{p*}, \quad \forall\, p \in \mathcal{P},\ \forall\, i \in \mathcal{N}, \tag{5a}$$

$$\sum_{i \in \mathcal{N}} \mu_i^{p*} = 0, \qquad \forall\, p \in \mathcal{P}, \tag{5b}$$

where $\left(\mu_i^{p*}\right)_{p \in \mathcal{P}, i \in \mathcal{N}}$ are the (unique) optimal dual variables. The dual variables arise in this unconstrained optimization problem because of the standard technique rewriting the public goods problem (4) as a private goods problem with equality constraints,

$$\left(x_1^*, x_2^*, \ldots, x_N^*\right) = \underset{x_1, \ldots, x_N \in \mathbb{R}^P}{\operatorname{argmax}} \sum_{i \in \mathcal{N}} v_i(x_i) \tag{6a}$$

$$\text{s.t.} \quad x_1 = x_2,\ x_2 = x_3, \cdots, x_N = x_1. \tag{6b}$$

The model above, leading up to (4), captures the possibility that one of the agents in $\mathcal{N}$ is a seller and thus his/her utility $v_i(x) = -c_i(x)$ is the negative of the cost of producing quantity x of the public good. In this case, a convex cost of production leads to a concave utility function. In general, for the public goods problem one can also explicitly assume a seller in the system who produces the quantity x and for whom the cost of production is a known (convex) function. In this case, the social welfare maximizing allocation contains the utility of the seller as well. For clarity, the seller is not considered in this model but if needed, this can accommodated in a straightforward manner.

3 A PUBLIC GOODS MECHANISM

A one-shot mechanism is defined by the triplet,

$$\left(\mathcal{M} = \mathcal{M}_1 \times \cdots \times \mathcal{M}_N, \left(\widehat{x}_1(\cdot), \ldots, \widehat{x}_N(\cdot)\right), \left(\widehat{t}_1(\cdot), \ldots, \widehat{t}_N(\cdot)\right)\right) \tag{7}$$

which consists of, for each agent $i \in \mathcal{N}$, the message space $\mathcal{M}_i$, the allocation function $\widehat{x}_i : \mathcal{M} \to \mathbb{R}^P$ and the tax function $\widehat{t}_i : \mathcal{M} \to \mathbb{R}$. Given a mechanism, a game $\mathfrak{G}$ is setup between the agents in $\mathcal{N}$, with action space $\mathcal{M}$ and utilities

$$u_i(m) = v_i(\widehat{x}_i(m)) - \widehat{t}_i(m). \tag{8}$$

The mechanism is said to *fully implement* the efficient allocation if

$$\widehat{x}_1(\widetilde{m}) = \widehat{x}_2(\widetilde{m}) = \cdots = \widehat{x}_N(\widetilde{m}) = x^* \tag{9}$$

for all pure strategy NE $\widetilde{m}$, where x^* is the efficient allocation from (4). Furthermore, the mechanism is said to be *budget balanced* at NE if $\sum_{i \in \mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0$ for all pure strategy NE $\widetilde{m}$. Finally, we call the mechanism *distributed* if for any agent $i \in \mathcal{N}$, the allocation $\widehat{x}_i(\cdot)$ and tax $\widehat{t}_i(\cdot)$ functions, instead of depending on the entire message $m = (m_j)_{j \in \mathcal{N}}$, depend only on m_i and $\left(m_j\right)_{j \in \mathcal{N}(i)}$, i.e., agent i and his/her immediate neighbors. Here $\mathcal{N}(i)$ are all the "out"-neighbors of i i.e., there exists an edge from i to j in graph $\mathcal{G}$ iff $j \in \mathcal{N}(i)$.

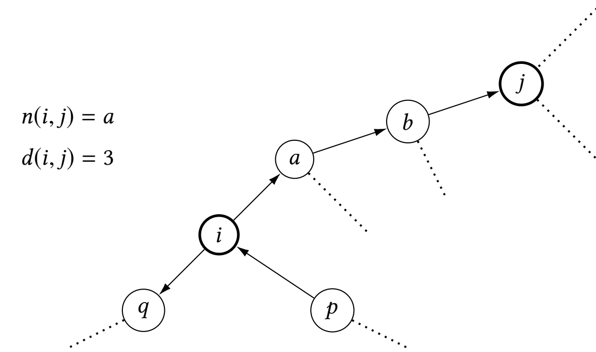

Figure 1: $n(i,j)$ and $d(i,j)$ for the strongly connected directed graph $\mathcal{G}$.

3.1 Mechanism

The presented mechanism below is for the special case of a single feature, i.e., $P = 1$, as it captures the essence of design. A natural extension to the general case is discussed at the end of this section.

Since the underlying graph $\mathcal{G}$ is strongly connected, for any pair of vertexes $i, j \in \mathcal{N}$, the following two quantities are well-defined. $d(i,j)$ is the length of the shortest path from i to j and $n(i,j) \in \mathcal{N}(i)$ is the out-neighbor of i such that the shortest path from i to j goes through the neighboring node $n(i,j)$. The two quantities are depicted in Fig. 1.

For any agent $i \in \mathcal{N}$, the message space is $\mathcal{M}_i = \mathbb{R}^{N+1}$. The message $m_i = (y_i, q_i)$ consists of agent i's contribution $y_i \in \mathbb{R}$ to the common public good and a surrogate/proxy $q_i = (q_i^1, \ldots, q_i^N) \in \mathbb{R}^N$ for the contributions of all the agents (including himself/herself).

The allocation function is defined as

$$\widehat{x}_i(m) = \frac{1}{N}\left(y_i + \sum_{r \in \mathcal{N}(i)} \frac{q_r^r}{\xi} + \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}}\right), \quad \forall\ i \in \mathcal{N}. \tag{10}$$

The tax function is

$$\widehat{t}_i(m) = \widehat{p}_i(m_{-i})\widehat{x}_i(m) + \left(q_i^i - \xi y_i\right)^2 + \sum_{r \in \mathcal{N}(i)} \left(q_i^r - \xi y_r\right)^2$$
$$+ \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \left(q_i^r - \xi q_{n(i,r)}^r\right)^2 + \frac{\delta}{2}\left(q_{n(i,i)}^i - \xi y_i\right)^2, \tag{11a}$$

$$\widehat{p}_i(m_{-i}) = \delta(N-1)\left(\frac{q_{n(i,i)}^i}{\xi} - \frac{1}{N-1}\left(\sum_{r \in \mathcal{N}(i)} \frac{q_r^r}{\xi} + \sum_{\substack{r \notin \mathcal{N}(i) \\ r \neq i}} \frac{q_{n(i,r)}^r}{\xi^{d(i,r)-1}}\right)\right),$$
$$\forall\ i \in \mathcal{N}, \tag{11b}$$

where $n(i,i) \in \mathcal{N}(i)$ is an arbitrarily chosen neighbor of i and $\xi \in (0,1)$, $\delta > 0$ are appropriately chosen parameters and their selection is discussed in Section 4 on Learning Guarantees.

The quantities, $n(\cdot,\cdot)$ and $d(\cdot,\cdot)$, are based on the graph $\mathcal{G}$. The only property of relevance here is that the two are related recursively i.e., $d(n(i,r),r) = d(i,r) - 1$. Thus if a designer wishes to avoid calculating the shortest path (possibly due to the high complexity) then n, d can be replaced by any valid neighbor and distance mapping, as long as they are related recursively as above.

3.2 Results

Fact 1 (Distributed). *The mechanism defined in* (10) *and* (11) *is distributed.*

This simply follows from the definition of $\widehat{x}_i(\cdot)$ and $\widehat{t}_i(\cdot)$ above, where only m_i and $\left(m_j\right)_{j\in\mathcal{N}(i)}$ are used.

Generally, for a public goods problem one expects the allocation function of the form $\widehat{x} : \mathcal{M} \to \mathbb{R}$ instead of $\left(\widehat{x}_i : \mathcal{M} \to \mathbb{R}\right)_{i\in\mathcal{N}}$ i.e, a single common allocation instead of N different allocation functions, one for each agent. However, owing to the informational constraints of the model, there does not exist[4] a single function $\widehat{x} : \mathcal{M} \to \mathbb{R}$ such that it depends only on m_i and $\left(m_j\right)_{j\in\mathcal{N}(i)}$, for every $i \in \mathcal{N}$. Since in general there are N different neighborhoods (one for each agent), the informational constraints necessitate the use of N allocation functions $\left(\widehat{x}_i\right)_{i\in\mathcal{N}}$. The multiple allocation functions are consistent with the model in (6). As each $\widehat{x}_i$ represents the level of the same public good, at NE we must have $\widehat{x}_i(\widetilde{m}) = \widehat{x}_j(\widetilde{m})$ for all $i, j \in \mathcal{N}$. The fact that $\widehat{x}_i(m) \neq \widehat{x}_j(m)$ for all $m \in \mathcal{M}$ means that the allocation is not feasible off-equilibrium.

One interpretation for N allocation functions is as follows: at the time of signing the mechanism contract, agents are aware of the informational constraints that the mechanism and its implied messaging must satisfy. Thus, each agent i can only use messages $\left(m_j\right)_{j\in\mathcal{N}(i)}$ to calculate their own best-response. The mechanism designer thus designs allocation function $\widehat{x}_i$ for agent i such that appropriate proxies[5] can be used for messages $\left(m_j\right)_{j\notin\mathcal{N}(i), j\neq i}$ of agents not in the neighborhood of i. Effectively thus, the mechanism designer accounts for the informational constraints through the appropriate design of proxies. A second interpretation is through the simulated learning process of NE. In this each agent "acts virtually" during the learning phase i.e., each agent takes only virtual actions during the learning phase, updating his/her action from one round to the next by observing only his/her neighbors message. Once the stationary point of the learning process is reached (NE), the "real action" is taken only once. The virtual interpretation is also used in most of the distributed optimization literature [3, 8, 17, 19].

The specific design of allocation and taxes in (10) and (11) is discussed below. The optimality conditions in (5) require that agents have global consensus on two aspects: allocation must be the same for all and sum of prices should be equal to zero. This is not straightforward as the agents are restricted to communicate with only their neighbors. Thus we introduce surrogate variables $q_i = (q_i^1, \ldots, q_i^N)$ which are known locally to agent i but at NE are expected to be proportional to the global demand $(y_1, \ldots, y_N)$. We design the second, third and fourth tax terms in (11a) such that agents are incentivized to duplicate global demand y onto locally available variables q_i.

To motivate the allocation function and the remaining part of the tax function consider the case of $\xi = 1$ and assume that the demands are duplicated on q_r as described above, i.e., $q_r = y$, $\forall\, r \in \mathcal{N}$. In this case all the factors involving ξ in (10) and (11) are 1. The allocation function $\widehat{x}_i(m)$ depends on $y_i, (q_r)_{r\in\mathcal{N}(i)}$ and is designed such that after taking into account the duplication it is proportional to $\sum_j y_j$. This facilitates the first consensus - all agents' allocation must be the same. The price $\widehat{p}_i(m_{-i})$ is designed such that it depends only on $(q_r)_{r\in\mathcal{N}(i)}$ and after taking into account the duplication it is proportional to $y_i - \frac{1}{N-1}\sum_{j\neq i} y_j$. This facilitates the second consensus - sum of prices over all agents is zero. With the above design principles, all the results of this section follow. We then introduce an additional fifth term in the tax, (11a), just for the purpose of achieving contraction in Section 4 (see proof of Theorem 4.3). Incentives provided by this term are in line with those already provided to the neighboring agent $n(i, i)$ through his/her third tax term, hence it doesn't interfere with the equilibrium results in this section. Finally, we set $\xi < 1$ so that the game $\mathfrak{G}$ can be contractive and adjust everything in the allocation and tax function correspondingly.

Lemma 3.1 (Concavity). *For any $i \in \mathcal{N}$ and $m_{-i} \in \mathcal{M}_{-i}$, utility $u_i(m)$ for the game $\mathfrak{G}$ is strictly concave in m_i. Thus, the best-response of agent i,*

$$\beta_i(m_{-i}) = \left(\tilde{y}_i(m_{-i}), \tilde{q}_i(m_{-i})\right) \triangleq \underset{m_i\in\mathcal{M}_i}{\operatorname{argmax}}\; u_i(m), \tag{12}$$

is unique (i.e., a single-valued function) and is defined by the first order conditions.

Proof. Please see Appendix A. ∎

Concavity of $u_i(m)$ follows from quadratic tax function, linear allocation function and the fact that $\widehat{p}_i$ doesn't depend on m_i. We verify concavity by showing the Hessian of $u_i(m)$ w.r.t. m_i to be negative definite. Second tax term in (11a) is only source of cross derivatives within components of $m_i = (y_i, q_i)$.

Theorem 3.2 (Full Implementation and Budget Balance). *For the game $\mathfrak{G}$, there exists a unique Nash equilibrium, $\widetilde{m} \in \mathcal{M}$, and the allocation at Nash equilibrium is efficient, i.e., $\widehat{x}(\widetilde{m}) = x^*$. Further, the total tax paid at Nash equilibrium $\widetilde{m}$ is zero, i.e., $\sum_{i\in\mathcal{N}} \widehat{t}_i(\widetilde{m}) = 0$.*

Proof. Please see Appendix B. ∎

As the optimality conditions (5) are sufficient, we show that at any NE the allocation and price satisfy (5). Thus if pure strategy NE exists, the corresponding allocation must be efficient. Existence and uniqueness is then established by providing a one-to-one map between NE and (x^*, μ^*).

3.2.1 Generalizing to multiple features ($P > 1$). For the general case, extend the presented mechanism by first increasing the message space such that for each agent $m_i = (m_i^p)_{p\in\mathcal{P}}$, i.e., each agent quotes demand and proxy $m_i^p = (y_i^p, q_i^p) \in \mathbb{R}^{N+1}$ separately for each good $p \in \mathcal{P}$. The allocation in this case is P–dimensional and the expression for $\widehat{x}_i^p(m)$ is the same as in the presented mechanism with $y_i, \left(q_r\right)_{r\in\mathcal{N}(i)}$ replaced by $y_i^p, \left(q_r^p\right)_{r\in\mathcal{N}(i)}$. The tax function is

$$\widehat{t}_i(m) = \sum_{p\in\mathcal{P}} \widehat{t}_i^{\,p}(m), \tag{13}$$

where the expression for $\widehat{t}_i^{\,p}$ is the same as in the presented mechanism, replacing m on RHS by $(m_i^p)_{i\in\mathcal{N}}$.

The proof of concavity is still completed by verifying that the Hessian of $u_i(m)$ w.r.t. m_i is negative definite. Calculations are a little different from the proof of Lemma 3.1 as the Hessian of $v_i(\cdot)$ in general has non-zero off-diagonal terms which means that the Hessian of u_i w.r.t. m_i has additional non-zero off-diagonal terms.

[4] apart from the trivial constant function, which cannot give a full implementation result.

[5] described in the paragraph below.

The Full implementation and Budget Balance results follow essentially from similar techniques as in the proof of Theorem 3.2.

4 DYNAMIC LEARNING GUARANTEES

This section provides the result for guaranteed convergence for a class of learning dynamics, when the mechanism defined in Section 3 is played repeatedly. As discussed in the Introduction, this makes the NE more stable w.r.t. information available to agents, and thus makes the mechanism ready for practical applications.

A learning dynamic is represented by $(\mu_n)_{n\geq 1} \subseteq \underset{i\in\mathcal{N}}{\times} \Delta(\mathcal{M}_i)$, where μ_n is a mixed strategy profile with product structure to be used at time n. Denote by $S(\mu_n)$ the support of the mixed strategy profile μ_n and denote by $m_n \in S(\mu_n)$ the realized action. Healy and Mathevet in [10] define the adaptive best-response (ABR) dynamics class of learnixng dynamics by restricting the support $S(\mu_n)$ in terms of past observed actions. Define the history $H_{n',n}$ as the set of observed actions between rounds n' and $n-1$. Define $B(\mathcal{M}')$ as the smallest closed ball centered at NE $\widehat{m}$ that contains the set $\mathcal{M}'$ (with any valid metric d on space $\mathcal{M}$). A learning dynamic is in the ABR class if any point in the support of the action at time n is no further from the NE than the best-response to any action that is no further from NE than the "worst-case" action that has been observed in some finite past.

Definition 4.1 (Adaptive Best-Response Learning Class [10]). A learning dynamic is an adaptive best-response dynamic if,

$$\forall\, n',\ \exists\, \widehat{n} > n',\ \text{s.t.}\ \forall\, n \geq \widehat{n}, \qquad S(\mu_n) \subseteq B\left(\beta\left(B\left(H_{n',n}\right)\right)\right), \quad (14)$$

where $\beta : \mathcal{M} \to \mathcal{M}$ is the best-response of the game.

The above is satisfied for instance if every agent puts belief zero over actions further from NE than the ones that he/she has observed in the past. Some well-known learning dynamics in the ABR class are as follows. Cournot best-response is the dynamics where at every time agents best-respond to the last round's action. More generally, $k-$period best-response is defined as the learning dynamic where at any time n, an agent i's strategy is a best-response to the mixed strategy of agents $j \neq i$ which are created using the observed empirical distribution from the actions of the previous $k-$rounds i.e., $\{m_{j,n-k}, \ldots, m_{j,n-1}\}$. In fact, the generalization of this is also in the ABR class, where at each time n, an agent i's strategy is the best-response to the mixed strategy of agents $j \neq i$ that is formed by taking any convex combination of the empirical distributions of actions observed in the previous $k-$rounds. Finally, fictitious Play [5, 9], which maintains empirical distribution of all the past actions (instead of k most recent ones) is also in ABR. The additional requirement for this is that the utility in the game should be strictly concave, which is true for the presented mechanisms (Lemma 3.1).

Definition 4.2 (Contractive Mechanism). A mechanism is called contractive if for any possible profile of utility function $\left(v_i(\cdot)\right)_{i\in\mathcal{N}}$, the induced game $\mathfrak{G}$ has a single-valued best-response function $\beta : \mathcal{M} \to \mathcal{M}$ that is a $d-$contraction mapping[6] (for any metric d on space $\mathcal{M}$).

[6] In this paper, we use the fact that h is a contraction mapping if the Jacobian has norm less than one, i.e., $\|\nabla h\| < 1$, where any matrix norm can be considered. Specifically, we consider the *row-sum* norm.

For the game induced by a contractive mechanism, by definition, there is a unique NE (due to the Banach fixed-point theorem). As shown in the previous section, the game $\mathfrak{G}$ indeed has a unique NE.

Fact 2 ([10, Theorem 1]). *If a game is contractive, then all ABR dynamics converge to the unique Nash equilibrium.*

It is shown below that the presented mechanism is contractive and thus owing to the above result there is guaranteed convergence for all learning dynamics in the ABR class. Contraction ensures convergence for the ABR class; this result is in the same vein as the one in the seminal work [14]. Milgrom and Roberts show that Supermodularity ensures convergence for the Adaptive Dynamics class of learning algorithms (also defined in [14]). Supermodularity requires that the best-response of any agent i is non-decreasing in the message m_j of any other agent $j \neq i$. The aim in this paper is to get guarantees for the ABR class, however it is shown below that the game $\mathfrak{G}$ is also supermodular and thus has guaranteed convergence for the Adaptive Dynamics class as well. As contraction is a more stringent condition than supermodularity, the ABR class is broader than the adaptive dynamics class.

Theorem 4.3 (Contraction). *The game $\mathfrak{G}$ defined in Section 3 is contractive and thus, all learning dynamics within the ABR dynamics class converge to the unique Nash equilibrium. Additionally, the game is also supermodular.*

Proof. Please see Appendix C. ∎

We begin by considering parameters ξ, δ such that $\xi \in (0,1)$ and $\delta > 0$. In order to get contraction, a further restriction $\xi \in (\sqrt{(N-1)/N}, 1)$ is imposed. This also makes the game supermodular. Finally, in order to accommodate any value of $\eta > 1$, the final selection of ξ requires it to be chosen very close to 1 and for δ to be selected as a function of ξ.

The proof of above relies on bounding (appropriately) the derivative of the inverse of $v_i'(\cdot)$. Generalizing to $P > 1$, the proof works by inverting the gradient ∇v_i and bounding its derivative (matrix). Bounds on the Hessian, (1), are then used to bound this derivative. From here onwards the proof steps follow analogously to the ones from the special case of $P = 1$.

4.1 Numerical Analysis of convergence

For numerical analysis we consider one feature ($P = 1$) and $N = 31$, $\eta = 25$. The agents' utility function as quadratic[7], $v_i(x) = \theta_i x^2 + \sigma_i x$. We have $v_i'' = 2\theta_i$ and thus the value for θ_i is chosen uniformly randomly in the range $(-\frac{\eta}{2}, -\frac{1}{2\eta})$. As the model doesn't impose any restriction on the first derivative, the value for σ_i is chosen uniformly randomly in the range $(10, 20)$.

From the proof of Theorem 4.3, one can numerically calculate the value of parameters ξ, δ. For the particular instance of the random θ, σ generated to be used for the plot below, the parameter values are: $\xi = 1 - 2.515 \times 10^{-4}$, $\delta = 0.9505$ when graph $\mathcal{G}$ is a full binary tree and $\xi = 1 - 10^{-3}$, $\delta = 0.8744$ when graph $\mathcal{G}$ is a sample of the Erdős-Reńyi random graph with only one connected component, where any two edges are connected with probability $p = 0.3$. The

[7] An example of quadratic utility function can be found in [20], for the model of demand side management in smart-grids.

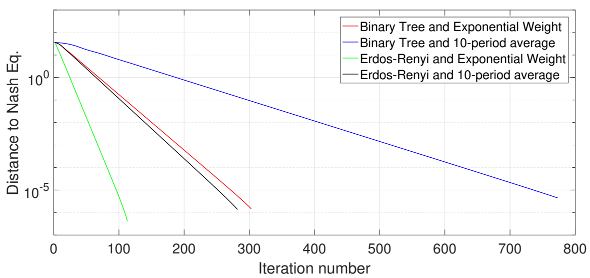

Figure 2: $\|m_n - \widetilde{m}\|_2$ vs. n for the public goods mechanism.

first represents a case of small average degree whereas the second represents the case of large average degree.

As the best-response is a contraction mapping, it is expected that any learning strategy that best-responds to some convex combination of (finite) past actions, converges at an exponential rate. Indeed, this is exactly observed from Fig. 2, where the absolute distance to NE is plotted. We consider two different learning dynamics. One where agents, at each time n, best-respond to an exponentially weighed average of past actions, i.e., $m_n = \beta\left(\frac{m_{n-1}}{2} + \frac{r_{n-1}}{2}\right)$ and $r_n = \frac{m_n}{2} + \frac{r_{n-1}}{2}$. The second learning dynamic is where agents best-respond to the arithmetic mean of past 10 rounds.

It is also evident from Fig. 2 that convergence is faster for more connected Erdős-Reńyi graph. This is expected since learning iterations facilitate information exchange, which is faster if there is more connectivity between agents. In fact, for both learning dynamics the convergence for the Erdős-Reńyi random graph is faster than either learning dynamic for Binary Tree. Comparing the two learning dynamics among themselves, we observe that the more aggressive exponential weighing leads to faster convergence compared to the learning dynamic that puts equal weight on each of the previous 10 actions. Finally, for Fig. 2, in each case the relative distance to NE, defined as $\|m_n - \widetilde{m}\|_2 / \|\widetilde{m}\|_1$, is in the order of 10^{-8} when the absolute distance to NE is 10^{-5}.

5 CONCLUSION

We consider the classical public goods (Lindahl) resource allocation problem for a distributed set of agents with informational constraints. We present a distributed mechanism which respects the informational constraints of the underlying graph. While for models with non-strategic agents, extensive research has been done in the field of distributed learning and optimization, this is not the case with mechanism design where agents are fully strategic.

The defined mechanism achieves Full implementation and Budget Balance i.e., for every possible profile $(v_i(\cdot))_{i\in\mathcal{N}}$ of utility functions, the induced game is shown to have a unique NE where the allocation at NE is efficient and taxes are budget balanced. Then we establish dynamic stability of the NE w.r.t. information available to agents by showing that the best-response of the induced game is a contraction mapping. This gives, using a result from [10], that every learning dynamic within the ABR dynamics class converges to the unique and efficient NE when the game is played repeatedly. The ABR class contains learning dynamics such as Cournot best-response, $k-$period best-response and fictitious Play.

Future Work. A scalable mechanism is desirable, where on average each agent's message space is of dimension $o(N)$. However, such an attempt might possibly require restrictions on either the underlying graph $\mathcal{G}$ or the upper bound on η that is admissible under the model. For the presented mechanism, the only restriction on the graph is that it is strongly connected and there is no upper bound on η. Another interesting direction to investigate is the possibility of extending the ABR class by making the convergence criteria broader. With contractive best-response, *almost sure* convergence is guaranteed within ABR class, but for instance if one is interested only in convergence in probability then a larger class of learning dynamics can be considered.

REFERENCES

[1] Timothy Besley and Stephen Coate. 2003. Centralized versus decentralized provision of local public goods: a political economy approach. *Journal of public economics* 87, 12 (2003), 2611–2637.

[2] Tilman Börgers and Rajiv Sarin. 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1 (1997), 1–14.

[3] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* 3, 1 (2011), 1–122.

[4] George W Brown. 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13, 1 (1951), 374–376.

[5] George W Brown. 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13, 1 (1951), 374–376.

[6] Yan Chen. 2002. A family of supermodular Nash mechanisms implementing Lindahl allocations. *Economic Theory* 19, 4 (2002), 773–790.

[7] Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. 2015. Near-optimal no-regret algorithms for zero-sum games. *Games and Economic Behavior* 92 (2015), 327–348.

[8] John C Duchi, Alekh Agarwal, and Martin J Wainwright. 2012. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic control* 57, 3 (2012), 592–606.

[9] Drew Fudenberg. 1998. *The theory of learning in games*. Vol. 2. MIT press.

[10] Paul J Healy and Laurent Mathevet. 2012. Designing stable mechanisms for economic environments. *Theoretical Economics* 7, 3 (2012), 609–661.

[11] Josef Hofbauer and William H Sandholm. 2002. On the global convergence of stochastic fictitious play. *Econometrica* 70, 6 (2002), 2265–2294.

[12] Leonid Hurwicz. 1979. Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points. *The Review of Economic Studies* 46, 2 (1979), 217–225.

[13] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, and others. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.

[14] Paul Milgrom and John Roberts. 1990. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica: Journal of the Econometric Society* (1990), 1255–1277.

[15] Dov Monderer and Lloyd S Shapley. 1996. Fictitious play property for games with identical interests. *Journal of economic theory* 68, 1 (1996), 258–265.

[16] Jonh F. Nash. 1950. *Non-cooperative Games*. Ph.D. Dissertation. Princeton University.

[17] Angelia Nedić and Asuman Ozdaglar. 2009. Distributed subgradient methods for multi-agent optimization. *IEEE Trans. Automat. Control* 54, 1 (2009), 48–61.

[18] Martin J Osborne and Ariel Rubinstein. 1994. *A course in game theory*. MIT press.

[19] Michael Rabbat and Robert Nowak. 2004. Distributed optimization in sensor networks. In *Proceedings of the 3rd international symposium on Information processing in sensor networks*. ACM, 20–27.

[20] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W. S. Won. 2012. Advanced Demand Side Management for the Future Smart Grid Using Mechanism Design. *IEEE Transactions on Smart Grid* 3, 3 (Sept 2012), 1170–1180.

[21] Shrutivandana Sharma and Demosthenis Teneketzis. 2012. Local public good provisioning in networks: A Nash implementation mechanism. *IEEE Journal on Selected Areas in Communication* 30, 11 (2012), 2105–2116.

[22] H Peyton Young. 2004. *Strategic learning and its limits*. Oxford University Press.

Appendix

A Distributed Mechanism for Public goods Allocation with Dynamic Learning Guarantees

Abhinav Sinha and Achilleas Anastasopoulos

A PROOF OF LEMMA 3.1 (CONCAVITY)

Proof. Since the allocation and tax functions are smooth and $v_i(\cdot)$ is continuously double-differentiable, to establish concavity we show that the Hessian of $u_i(m)$ w.r.t. m_i is negative definite i.e., $H \prec 0$. Once this is established, the optimization for best-response

$$\beta_i(m_{-i}) = \left(\tilde{y}_i(m_{-i}), \tilde{q}_i(m_{-i})\right) \triangleq \underset{m_i \in \mathcal{M}_i}{\text{argmax}}\ u_i(m). \tag{1}$$

has a strictly concave objective and an unbounded constraint set. Thus it has a unique maximizer, defined by the first order derivative conditions.

The Hessian is of size $(N+1) \times (N+1)$ and we have

$$H_{11} = \frac{\partial^2 u_i(m)}{\partial y_i^2} = \frac{v_i''(\widehat{x}_i(m))}{N^2} - (2+\delta)\xi^2, \tag{2a}$$

$$H_{(j+1)1} = H_{1(j+1)} = \frac{\partial^2 u_i(m)}{\partial y_i \partial q_i^j} = \begin{cases} 0 & \text{for } j \in \mathcal{N},\ j \neq i, \\ 2\xi & \text{for } j = i, \end{cases} \tag{2b}$$

$$H_{(j+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial (q_i^j)^2} = -2, \quad \forall\, j \in \mathcal{N}, \tag{2c}$$

$$H_{(r+1)(j+1)} = \frac{\partial^2 u_i(m)}{\partial q_i^r \partial q_i^j} = 0, \quad \forall\, j, r \in \mathcal{N},\ j \neq r. \tag{2d}$$

The characteristic equation, $\text{Det}\,(H - xI) = 0$, becomes

$$(x+2)^{N-1}\left((x+2)(x - H_{11}) - 4\xi^2\right) = 0. \tag{3}$$

This implies that $N-1$ eigenvalues of H are -2 and the remaining two eigenvalues satisfy $x^2 + (2 - H_{11})x + 2\delta\xi^2 - \frac{2}{N^2} v_i''(\widehat{x}_i(m)) = 0$. Since H is a symmetric matrix, all its eigenvalues are real. Due to $v_i''(\cdot) < 0$, the product of roots in the above quadratic equation is positive and the sum of roots is negative. This gives that the remaining two eigenvalues of H are also negative. ■

B PROOF OF THEOREM 3.2 (FULL IMPLEMENTATION)

Proof. For the social welfare maximization problem,

$$x^* = \operatorname*{argmax}_{x \in \mathbb{R}} \sum_{i \in \mathcal{N}} v_i(x), \tag{4}$$

the optimality conditions

$$\frac{\mathrm{d}v_i(x^*)}{\mathrm{d}x} = \mu_i^*, \quad \forall\, i \in \mathcal{N}, \tag{5a}$$

$$\sum_{i \in \mathcal{N}} \mu_i^* = 0, \tag{5b}$$

are sufficient. Thus in order to prove that the corresponding allocation at Nash equilibrium is efficient, we show that at any Nash equilibrium $\overline{m} = (\overline{y}, \overline{q}) \in \mathcal{M}$, the allocation $\left(\widehat{x}_i(\overline{m})\right)_{i \in \mathcal{N}}$ and prices $\left(\widehat{p}_i(\overline{m})\right)_{i \in \mathcal{N}}$ satisfy the optimality conditions as x^* and μ^*, respectively. Then using an invertibility argument we show existence and uniqueness of Nash equilibrium.

Using Lemma 3.1, at any Nash equilibrium $\overline{m}$ we have: $\nabla_{m_i} u_i(\overline{m}) = 0, \forall\, i \in \mathcal{N}$. This gives

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial y_i} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial y_i} = 0, \quad \forall\, i \in \mathcal{N}, \tag{6a}$$

$$\frac{\partial v_i(\widehat{x}_i(\overline{m}))}{\partial q_i^r} - \frac{\partial \widehat{t}_i(\overline{m})}{\partial q_i^r} = 0, \quad \forall\, r \in \mathcal{N},\ i \in \mathcal{N}. \tag{6b}$$

Using the definition of allocation and tax function from the mechanism, this becomes, $\forall\, i \in \mathcal{N}$,

$$\frac{1}{N}\left(v_i'(\widehat{x}_i(\overline{m})) - \widehat{p}_i(\overline{m}_{-i})\right) + 2\xi(\overline{q}_i^i - \xi\overline{y}_i) + \delta\xi(\overline{q}_{n(i,i)}^i - \xi\overline{y}_i) = 0, \tag{7a}$$

$$\overline{q}_i^r = \begin{cases} \xi\overline{y}_i & \text{for } r = i, \\ \xi\overline{y}_r & \text{for } r \in \mathcal{N}(i), \\ \xi\overline{q}_{n(i,r)}^r & \text{for } r \notin \mathcal{N}(i) \text{ and } r \neq i, \end{cases} \quad . \tag{7b}$$

For any distinct pair of vertexes i, r, denote by $\{i, i_1, i_2, \ldots, i_{d(i,r)} = r\}$ the ordered vertexes in the shortest path between i and r, where $i_1 = n(i, r) \in \mathcal{N}(i)$. Since the shortest path between i and r contains the shortest path between i_k and r, for any $k < d(i, r)$, we have $n(i_k, r) = i_{k+1}$. Using the third sub-equation in (7b) repeatedly, replacing i by i_k gives,

$$\overline{q}_i^r = \xi\, \overline{q}_{i_1}^r = \xi^2\, \overline{q}_{i_2}^r = \cdots = \xi^{d(i,r)-1}\, \overline{q}_{i_{d(i,r)-1}}^r. \tag{8}$$

Now using the second sub-equation of (7b), replacing i by $i_{d(i,r)-1}$ and noting $r \in \mathcal{N}(i_{d(i,r)-1})$, gives $\overline{q}_{i_{d(i,r)-1}}^r = \xi\, \overline{y}_r$. This combined with the above equation gives that (7b) implies

$$\overline{q}_i^r = \begin{cases} \xi \overline{y}_i & \text{for } r = i, \\ \xi^{d(i,r)} \overline{y}_r & \text{for } r \neq i, \end{cases} \qquad \forall\, i \in \mathcal{N}. \tag{9}$$

Using the above and then combining (7a) with the definition of allocation and tax functions gives, $\forall\, i \in \mathcal{N}$,

$$v_i'(\widehat{x}_i(\overline{m})) = \widehat{p}_i(\overline{m}_{-i}), \tag{10a}$$

$$\widehat{x}_i(\overline{m}) = \frac{1}{N} \sum_{j \in \mathcal{N}} \overline{y}_j, \tag{10b}$$

$$\widehat{p}_i(\overline{m}_{-i}) = \delta(N-1) \left(\overline{y}_i - \frac{1}{N-1} \sum_{j \neq i} \overline{y}_j \right). \tag{10c}$$

(10b) implies $\widehat{x}_i(\overline{m}) = \widehat{x}_r(\overline{m})$ for any $i, r \in \mathcal{N}$ and (10c) gives $\sum_{r \in \mathcal{N}} \widehat{p}_i(\overline{m}_{-i}) = 0$. Thus, the allocation-price pair

$$\left(\frac{1}{N} \sum_{j \in \mathcal{N}} \overline{y}_j \,,\, \left(\delta(N-1) \left(\overline{y}_i - \frac{1}{N-1} \sum_{j \neq i} \overline{y}_j \right) \right)_{i \in \mathcal{N}} \right) \tag{11}$$

satisfy the optimality conditions, (5), as (x^*, μ^*). Since the optimality conditions are sufficient, the allocation at any Nash equilibrium $\overline{m}$ is the efficient allocation x^*.

For existence and uniqueness, consider the following set of linear equations that must be satisfied at any Nash equilibrium $\overline{m}$,

$$x^* = \frac{1}{N}\sum_{j\in\mathcal{N}} \overline{y}_j, \tag{12a}$$

$$\mu_i^* = \delta(N-1)\left(\overline{y}_i - \frac{1}{N-1}\sum_{j\neq i}\overline{y}_j\right), \quad \forall\, i \in \mathcal{N}. \tag{12b}$$

Here $\left(\overline{y}_j\right)_{j\in\mathcal{N}}$ are the variables and (x^*, μ^*) are fixed - since they are uniquely defined by the optimization, (4). The above equations can be inverted to give the unique solution as,

$$\overline{y}_i = x^* + \frac{\mu_i^*}{\delta N}, \quad \forall\, i \in \mathcal{N}. \tag{13}$$

Furthermore, using above and (9), the values for $\left(\overline{q}_i^r\right)_{i,r\in\mathcal{N}}$ can also be calculated uniquely. Since a solution for $\overline{m} = (\overline{y}, \overline{q})$ in terms of x^*, μ^* exists, existence of Nash equilibrium is guaranteed. Also, since this solution is unique, there is a unique Nash equilibrium.

For Budget Balance, we have the following. By the characterization from above we know that at Nash Equilibrium $\widetilde{m}$, all tax terms from the definition of $\widehat{t}_i(m)$, other than $\widehat{p}_i(\widetilde{m}_{-i})\widehat{x}_i(\widetilde{m})$, are zero. Furthermore, the prices are equal to μ_i^* and each allocation is equal to x^*. Thus,

$$\sum_{i\in\mathcal{N}} \widehat{t}_i(\widetilde{m}) = \sum_{i\in\mathcal{N}} \mu_i^* x^* = x^* \sum_{i\in\mathcal{N}} \mu_i^* = x^* \cdot 0 = 0, \tag{14}$$

since the optimal dual variables $\left(\mu_i^*\right)_{i\in\mathcal{N}}$ satisfy, (5b), $\sum_{i\in\mathcal{N}} \mu_i^* = 0$. ■

C PROOF OF THEOREM 4.3 (CONTRACTION)

Proof. The game is contractive if the matrix norm of the Jacobian of best-response $\beta = \left(\beta_i\right)_{i\in\mathcal{N}} = \left(\tilde{y}_i, \tilde{q}_i\right)_{i\in\mathcal{N}}$ is smaller than unity, i.e., $\|\nabla\beta\| < 1$. We use the row-sum norm for this, and in this proof verify specifically the following set of conditions,

$$\sum_{r\in\mathcal{N},\, r\neq i}\left(\left|\frac{\partial \tilde{y}_i}{\partial y_r}\right| + \sum_{j\in\mathcal{N}}\left|\frac{\partial \tilde{y}_i}{\partial q_r^j}\right|\right) < 1, \quad \forall\, i \in \mathcal{N}, \tag{15a}$$

$$\sum_{r\in\mathcal{N},\, r\neq i}\left(\left|\frac{\partial \tilde{q}_i^w}{\partial y_r}\right| + \sum_{j\in\mathcal{N}}\left|\frac{\partial \tilde{q}_i^w}{\partial q_r^j}\right|\right) < 1, \quad \forall\, w \in \mathcal{N},\ \forall\, i \in \mathcal{N}. \tag{15b}$$

Consider any agent $i \in \mathcal{N}$, for the best-response $\tilde{q}_i$ we have

$$\tilde{q}_i^w = \begin{cases} \xi \tilde{y}_i & \text{for } w = i, \\ \xi y_w & \text{for } w \in \mathcal{N}(i), \\ \xi q^w_{n(i,w)} & \text{for } w \notin \mathcal{N}(i) \text{ and } w \neq i. \end{cases} \tag{16}$$

Thus, by choosing $\xi \in (0,1)$, all conditions within (15b) are satisfied where $w \neq i$. Next, we verify conditions in (15a). Once this is done, then in conjunction with $\xi \in (0,1)$, the conditions from (15b) with $w = i$ are also automatically verified.

For the best-response $\tilde{y}_i$, we have the following relation

$$\frac{1}{N}\big(v_i'(\widehat{x}_i(m)) - \widehat{p}_i(m_{-i})\big) + 2\xi(\tilde{q}_i^i - \xi\tilde{y}_i) + \delta\xi(q^i_{n(i,i)} - \xi\tilde{y}_i) = 0, \tag{17a}$$

$$\Rightarrow \frac{1}{N}\big(v_i'(\widehat{x}_i(m)) - \widehat{p}_i(m_{-i})\big) + \delta\xi(q^i_{n(i,i)} - \xi\tilde{y}_i) = 0. \tag{17b}$$

In the above relation, $\widehat{x}_i(m)$ is evaluated at $\tilde{y}_i$ instead of y_i. Also, this relation implicitly defines $\tilde{y}_i$. Differentiating this equation w.r.t. $\big(q^r_{n(i,r)}\big)_{r\in\mathcal{N}}$ gives

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,i)}} - \frac{\delta(N-1)}{N\xi} + \delta\xi\left(1 - \xi\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,i)}}\right) = 0, \qquad r = i, \tag{18a}$$

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\left(\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,r)}} + \frac{1}{\xi}\right) + \frac{\delta}{N\xi} - \delta\xi^2\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,r)}} = 0, \quad \forall\, r \in \mathcal{N}(i), \tag{18b}$$

$$\frac{v_i''(\widehat{x}_i(m))}{N^2}\left(\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,r)}} + \frac{1}{\xi^{d(i,r)-1}}\right) + \frac{\delta}{N\xi^{d(i,r)-1}} - \delta\xi^2\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,r)}} = 0, \quad \forall\, r \notin \mathcal{N}(i),\ r \neq i, \tag{18c}$$

which implies

$$\frac{\partial\tilde{y}_i}{\partial q^r_{n(i,r)}} = \frac{1}{\dfrac{v_i''(\widehat{x}_i(m))}{N^2} - \delta\xi^2} \times \begin{cases} \dfrac{\delta(N-1)}{N\xi} - \delta\xi & \text{for } r = i, \\ -\dfrac{\delta}{N\xi} - \dfrac{v_i''(\widehat{x}_i(m))}{N^2\xi} & \text{for } r \in \mathcal{N}(i), \\ -\dfrac{\delta}{N\xi^{d(i,r)-1}} - \dfrac{v_i''(\widehat{x}_i(m))}{N^2\xi^{d(i,r)-1}} & \text{for } r \notin \mathcal{N}(i),\ r \neq i. \end{cases} \tag{19}$$

In the notation used above, for any $r \in \mathcal{N}(i)$, $n(i,r) = r$. All other partial derivative of $\tilde{y}_i$ are zero. With all this condition in (15a) becomes,

$$\left|\frac{\delta(N-1)}{N\xi} - \delta\xi\right| + \left|-\frac{\delta}{N} - \frac{v_i''(\widehat{x}_i(m))}{N^2}\right|\left(\sum_{r\in\mathcal{N}(i)} \frac{1}{\xi}\right) + \left|-\frac{\delta}{N} - \frac{v_i''(\widehat{x}_i(m))}{N^2}\right|\left(\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{\xi^{d(i,r)-1}}\right) < \delta\xi^2 - \frac{v_i''(\widehat{x}_i(m))}{N^2}. \tag{20}$$

We impose the condition

$$\xi \in \left(\sqrt{\frac{N-1}{N}}, 1\right) \tag{21}$$

so that the expression inside the first absolute value term in above is negative. To simplify the other expressions containing absolute value, we utilize the lower bound from the assumption $v_i''(\cdot) \in (-\eta, -\frac{1}{\eta})$. Set

$$\eta < N\delta, \tag{22}$$

so that the remaining expressions inside absolute value in (20) are guaranteed to be negative. With this, (20) becomes

$$\left[-v_i''(\widehat{x}_i(m))\right]\left(1 + \sum_{r\in\mathcal{N}(i)} \frac{1}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{\xi^{d(i,r)-1}}\right) > N\delta\left(\frac{1}{\xi}\left[\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}} \frac{1}{\xi^{d(i,r)-2}} - (N - |\mathcal{N}(i)| - 1)\right] + N\xi(1-\xi)\right), \tag{23}$$

where $N - |\mathcal{N}(i)| - 1$ is the number of agents in the system except agent i and all his/her neighbors in $\mathcal{N}(i)$. Clearly the LHS above is positive. For any $r \in \mathcal{N}(i)$ and $r \neq i$, we have $d(i,r) \geq 2$. On the RHS, inside the square brackets there are exactly $N - |\mathcal{N}(i)| - 1$ terms in the summation and each term is of the form ξ^{-k}, for some $k \geq 0$. Since $\xi < 1$, this gives that even the RHS is positive. Utilizing the upper bound from $v_i''(\cdot) \in (-\eta, -\frac{1}{\eta})$, a sufficient condition to verify (23) is

$$\eta < \frac{1}{N\delta}\frac{C_i}{D_i} \tag{24}$$

where C_i, D_i are the expression inside the curved bracket on the LHS and RHS of (23), respectively. Combining this with the condition in (22), a sufficient condition for verifying (15a) is

$$\eta < \min\left(N\delta\,,\ \frac{1}{N\delta}\frac{C_i}{D_i}\right), \quad \forall\, i \in \mathcal{N}. \tag{25}$$

The proof is complete as long as η satisfies above. However, in our model we would like to accommodate any value of $\eta > 1$ and this requires selecting parameters ξ, δ appropriately. Set

$$\delta = \frac{1}{N}\sqrt{\min_{i\in\mathcal{N}}\left(\frac{C_i}{D_i}\right)}, \tag{26}$$

in above to get the sufficient condition as $\eta^2 < \min_{i\in\mathcal{N}}\left(\frac{C_i}{D_i}\right)$, i.e.,

$$\eta^2 < \min_{i\in\mathcal{N}}\left\{\left(1 + \sum_{r\in\mathcal{N}(i)}\frac{1}{\xi} + \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-1}}\right)\Bigg/\left(\frac{1}{\xi}\left[\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-2}} - (N - |\mathcal{N}(i)| - 1)\right] + N\xi(1-\xi)\right)\right\}. \tag{27}$$

We want to select ξ such that the RHS above can be made arbitrarily large, whilst satisfying (21). For this, first note that for any $i \in \mathcal{N}$ the numerator of the RHS is greater than 1, hence it is bounded away from zero. Secondly, the denominator can be made arbitrarily close to 0 by choosing ξ close enough to 1. This can be seen by rewriting

$$\sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\frac{1}{\xi^{d(i,r)-2}} - (N - |\mathcal{N}(i)| - 1) = \sum_{\substack{r\notin\mathcal{N}(i)\\ r\neq i}}\left(\frac{1}{\xi^{d(i,r)-2}} - 1\right) = \sum_{\substack{r\in\mathcal{N}\\ d(i,r)\geq 3}}\left(\frac{1}{\xi^{d(i,r)-2}} - 1\right), \tag{28a}$$

$$\Rightarrow\ D_i = \frac{1}{\xi}\sum_{\substack{r\in\mathcal{N}\\ d(i,r)\geq 3}}\left(\frac{1}{\xi^{d(i,r)-2}} - 1\right) + N\xi(1-\xi), \tag{28b}$$

where for any given $k \geq 1, \epsilon > 0$, choose $\xi \in \left(\left(\frac{1}{1+\epsilon}\right)^{1/k}, 1\right)$ to have $\left(\xi^{-k} - 1\right) < \epsilon$. Note that this is consistent with (21). The remaining term $N\xi(1-\xi)$ can also made made arbitrarily small by choosing ξ close enough to 1. Finally, it is clear from above that the denominator D_i can made arbitrarily small concurrently for all $i \in \mathcal{N}$.

This completes the proof of Contraction. The fact that the game is supermodular follows from the preceding analysis, where the parameters ξ, δ are chosen such that each expression in (19) is positive (and similarly all partial derivatives arising out of (16) are positive). Which implies that the best-response $\beta_i(m_{-i})$ is non-decreasing w.r.t message m_k for any other agent $k \neq i$. ■