

Multi-Scale Games: Representing and Solving Games on Networks with Group Structure

Kun Jin

University of Michigan, Ann Arbor
kunj@umich.edu

Yevgeniy Vorobeychik

Washington University in St. Louis
yvorobeychik@wustl.edu

Mingyan Liu

University of Michigan, Ann Arbor
mingyan@umich.edu

Abstract

Network games provide a natural machinery to compactly represent strategic interactions among agents whose payoffs exhibit sparsity in their dependence on the actions of others. Besides encoding interaction sparsity, however, real networks often exhibit a multi-scale structure, in which agents can be grouped into communities, those communities further grouped, and so on, and where interactions among such groups may also exhibit sparsity. We present a general model of multi-scale network games that encodes such multi-level structure. We then develop several algorithmic approaches that leverage this multi-scale structure, and derive sufficient conditions for convergence of these to a Nash equilibrium. Our numerical experiments demonstrate that the proposed approaches enable orders of magnitude improvements in scalability when computing Nash equilibria in such games. For example, we can solve previously intractable instances involving up to 1 million agents in under 15 minutes.

1 Introduction

Strategic interactions among interconnected agents are commonly modeled using the network, or graphical, game formalism (Kearns, Littman, and Singh, 2001; Jackson and Zenou, 2015). In such games, the utility of an agent depends on his own actions as well as those by its network neighbors. Many variations of games on networks have been considered, with applications including the provision of public goods (Allouch, 2015; Buckley and Croson, 2006; Khalili, Zhang, and Liu, 2019; Yu et al., 2020), security (Hota and Sundaram, 2018; La, 2016; Vorobeychik and Letchford, 2015), and financial markets (Acemoglu et al., 2012).

While network games are a powerful modeling framework, they fail to capture a common feature of human organization: groups and communities. Indeed, investigation of communities, or close-knit groups, in social networks is a major research thread in network science. Moreover, such groups often have a hierarchical structure (Clauset, Moor, and Newman, 2008; Girvan and Newman, 2002). For example, strategic interactions among organizations in a marketplace often boil down to interactions among their constituent business units, which are, in turn, comprised of individual decision makers. In the end, it is those lowest-




Figure 1: An illustration of a multi-scale (3-level) network.

level agents who ultimately accrue the consequences of these interactions (for example, corporate profits would ultimately benefit individual shareholders). Moreover, while there are clear interdependencies among organizations, individual utilities are determined by a combination of individual actions of some agents, together with *aggregate* decisions by the groups (e.g., business units, organizations). For example, an employee’s bonus is determined in part by their performance in relation to their co-workers, and in part by how well their employer (organization) performs against its competitors in the marketplace.

We propose a novel *multi-scale game model* that generalizes network games to capture such hierarchical organization of individuals into groups. Figure 1 offers a stylized example in which three groups (e.g., organizations) are comprised of 2-3 subgroups each (e.g., business units), which are in turn comprised of 2-5 individual agents. Specifically, our model includes an explicit hierarchical network structure that organizes agents into groups across a series of levels. Further, each group is associated with an action which deterministically aggregates the decisions by its constituent agents. The game is grounded at the lowest level, where the agents are associated with scalar actions and utility functions that have modular structure in the strategies taken at each level of the game. For example, in Figure 1, the utility function of an individual member a_j of level-3 group $a_3^{(3)}$ is a function of the strategies of (i) a_j ’s immediate neighbors

(represented by links between pairs of filled-in circles), (ii) a_j 's level-2 group and its network neighbor (the small hollow circles), and (iii) a_j 's level-3 group, $a_3^{(3)}$ (large hollow circle) and its network neighbors, $a_1^{(3)}$ and $a_2^{(3)}$.

Our next contribution is a series of iterative algorithms for computing pure strategy Nash equilibria that explicitly leverage the proposed multi-scale game representation. The first of these simply takes advantage of the compact game representation in computing equilibria. The second algorithm we propose offers a further innovation through an iterative procedure that alternates between game levels, treating groups themselves as pseudo-agents in the process. We present sufficient conditions for the convergence of this algorithm to a pure strategy Nash equilibrium through a connection to Structured Variational Inequalities (He, Yang, and Wang, 2000), although the result is limited to games with two levels. To address the latter limitation, we design a third iterative algorithm that now converges even in games with an arbitrary number of levels.

Our final contribution is an experimental evaluation of the proposed algorithms compared to best response dynamics. In particular, we demonstrate orders of magnitude improvements in scalability, enabling us to solve games that cannot be solved using a conventional network game representation.

Related Work: Network games have been an active area of research; see e.g., surveys by Jackson and Zenou (2015) and Bramoullé and Kranton (2016). We now review the most relevant papers. Conditions for the existence, uniqueness and stability of Nash equilibria in network games under general best responses are studied in (Parise and Ozdaglar, 2019; Naghizadeh and Liu, 2017; Scutari et al., 2014; Bramoullé, Kranton, and D'amours, 2014). Variational inequalities (VI) are used in these works to analyze the fixed point and contraction properties of the best response mappings. It is identified in Parise and Ozdaglar (2019); Naghizadeh and Liu (2017); Scutari et al. (2014) that when the Jacobian matrix of the best response mapping is a P-matrix or is positive definite, a feasible unique Nash equilibrium exists and can be obtained by best-response dynamics (Scutari et al., 2014; Parise and Ozdaglar, 2019). In this paper, we extended the analysis of equilibrium and best responses for a conventional network game to that in a multi-scale network game, where the utility functions are decomposed into separable utility components to which best responses are applied separately. This is similar to the generalization from a conventional VI problem to an SVI problem (He, Yang, and Wang, 2000; He, 2009; He and Yuan, 2012; Bnouchachem, Benazza, and Khalfaoui, 2013) problem.

Previous works on network games that involve group or community structure focus on finding such structures; e.g., community detection in networks using game theoretic methods have been studied in (Mcsweeney, Mehrotra, and Oh, 2017; Newman, 2004; Alvari, Hajibagheri, and Sukthankar, 2014). By contrast, our work focuses on analyzing a network game with a given group/community structure, and using the structure as an analytical tool for the analysis of equilibrium and best responses.

2 Preliminaries

A general *normal-form game* is defined by a set of agents (players) $I = \{1, \dots, N\}$, with each agent a_i having an action/strategy space K_i and a utility function $u_i(x_i, \mathbf{x}_{-i})$ that i aims to maximize; $x_i \in K_i$ and \mathbf{x}_{-i} denotes the actions by all agents other than i . We term the collection of strategies of all agents \mathbf{x} a *strategy profile*. We assume $K_i \subset \mathbb{R}$ is a compact set.

We focus on computing a *Nash equilibrium (NE)* of a normal-form game, which is a strategy profile with each agent maximizing their utility given the strategies of others. Formally, \mathbf{x}^* is a *Nash equilibrium* if for each agent i ,

$$x_i^* \in \operatorname{argmax}_{x_i \in K_i} u_i(x_i, \mathbf{x}_{-i}^*). \quad (1)$$

A *network game* encodes structure in the utility functions such that they only depend on the actions by network neighbors. Formally, a network game is defined over a weighted graph (I, E) , with each node an agent and E is the set of edges; the agent's utility $u_i(x_i, \mathbf{x}_{-i})$ reduces to $u_i(x_i, \mathbf{x}_{I_i})$, where I_i is the set of network neighbors of i , although we will frequently use the former for simplicity.

An agent's best response is its best strategy given the actions taken by all the other agents. Formally, the best response is a set defined by

$$BR_i(\mathbf{x}_{-i}, u_i) = \operatorname{argmax}_{x_i} u_i(x_i, \mathbf{x}_{-i}). \quad (2)$$

Whenever we deal with games that have a unique best response, we will use the singleton best response set to also refer to the player's best response strategy (the unique member of this set).

Clearly, a NE of a game is a fixed point of this best response correspondence. Consequently, one way to compute a NE of a game is through *best response dynamics (BRD)*, which is a process whereby agents iteratively and asynchronously (that is, one agent at a time) take the others' actions as fixed values and play a best response to them.

We are going to use this BRD algorithm as a major building block below. One important tool that is useful for analyzing BRD convergence is *Variational Inequalities (VI)*. To establish the connection between NE and VI we assume the utility functions $u_i, \forall i = 1, \dots, N$, are continuously twice differentiable. Let $K = \prod_{i=1}^N K_i$ and define $F : \mathbb{R}^N \rightarrow \mathbb{R}^N$ as follows:

$$F(\mathbf{x}) := \left(-\nabla_{x_i} u_i(\mathbf{x}) \right)_{i=1}^N. \quad (3)$$

Then \mathbf{x}^* is said to be a solution to $VI(K, F)$ if and only if

$$(\mathbf{x} - \mathbf{x}^*)^T F(\mathbf{x}^*) \geq 0, \forall \mathbf{x} \in K. \quad (4)$$

In other words, the solution set to $VI(K, F)$ is equivalent to the set of NE of the game. Now, we can define the condition that will guarantee the convergence of BRD.

Definition 1. The P_Υ condition: The Υ matrix generated from $F : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is given as follows

$$\Upsilon(F) = \begin{bmatrix} \alpha_1(F) & -\beta_{1,2}(F) & \cdots & -\beta_{1,N}(F) \\ -\beta_{2,1}(F) & \alpha_2(F) & \cdots & -\beta_{2,N}(F) \\ \vdots & \vdots & \ddots & \vdots \\ -\beta_{N,1}(F) & -\beta_{N,2}(F) & \cdots & \alpha_N(F) \end{bmatrix},$$

$\alpha_i(F) = \inf_{\mathbf{x} \in K} \|\nabla_i F_i\|_2$, $\beta_{i,j}(F) = \sup_{\mathbf{x} \in K} \|\nabla_j F_i\|_2$, $i \neq j$. If $\Upsilon(F)$ is a P -matrix, that is, if all of its principal components have a positive determinant, then we say F satisfies the P_Υ condition.

Theorem 1. (Scutari et al., 2014) *If F satisfies the P_Υ condition, then F is strongly monotone on K , and $VI(K, F)$ has a unique solution. Moreover, BRD converges to the unique NE from an arbitrary initial state.*

3 A Multi-Scale Game Model

Consider a conventional network (graphical) game with the set I of N agents situated on a network $G = (I, E)$, each with a utility function $u_i(x_i, \mathbf{x}_{I_i})$, with I_i the set of i 's neighbors, I the full set of agents/nodes and E the set of edges connecting them.¹ Suppose that this network G exhibits the following structure and feature of the strategic dependence among agents: agents can be partitioned into a collection of groups $\{S_k\}$, where k is a group index, and an agent a_i in the k th group (i.e., $a_i \in S_k$) has a utility function that depends (i) on the strategies of its network neighbors in S_k , and (ii) *only on the aggregate strategies* of groups other than k (see, e.g., Fig. 1). Further, these groups may go on to form larger groups, whose aggregate strategies impact each other's agents, giving rise to a *multi-scale* structure of the network. This kind of structure is very natural in a myriad of situations. For example, members of criminal organizations take stock of individual behavior by members of their own organization, but their interactions with other organizations (criminal or otherwise) are perceived in group terms (e.g., how much another group has harmed theirs). A similar multi-level interaction structure exists in national or ethnic conflicts, organizational competition in a market place, and politics. Indeed, a persistent finding in network science is that networks exhibit a multi-scale interaction structure (i.e., communities, and hierarchies of communities) (Girvan and Newman, 2002; Clauset, Moor, and Newman, 2008).

We present a general model to capture such multi-scale structure. Formally, an L -level structure is given by a hierarchical graph structure $\{G^{(l)}\}$ for each level l , $1 \leq l \leq L$, where $G^{(l)} = (\{S_k^{(l)}\}_k, E^{(l)})$ represents the level- l structure. The first component, $\{S_k^{(l)}\}_k$ prescribes a partition, where agents in level $l - 1$ form disjoint groups given by this partition; each group is viewed as an agent in level l , denoted as $a_k^{(l)}$. Notationally, while both $a_k^{(l)}$ and $S_k^{(l)}$ bear the superscript (l) , the former refers to a level- l agent, while the latter is the group (of level- $(l - 1)$ agents) that the former represents. The set of level- l agents is denoted by $I^{(l)}$ and their total number $N^{(l)}$. The second component, $E^{(l)}$, is a set of edges that connect level- l agents, encoding the dependence relationship among the groups they represent. This structure is anchored in level 1 (the lowest level), where sets $S_k^{(1)}$ are singletons, corresponding to agents a_k in the game, who constitute the set I .

To illustrate, the multi-scale structure shown in Fig. 1 is

¹The edges are generally weighted, resulting in a weighted adjacency matrix on which the utility depends.

given by $G^{(1)} = G = (\{S_k^{(1)}\}_k = I, E^{(1)} = E)$, as well as how level-1 agents are grouped into level-2 agents, how level-2 agents are further grouped into level-3 agents, and the edges connecting these groups at each level.

It should be obvious that the above multi-scale representation of a graphical game is a generalization of a conventional graphical game, as any such game essentially corresponds to a $L = 1$ multi-scale representation. On the other hand, not all conventional graphical games have a meaningful $L > 1$ multi-scale representation (with non-singleton groups of level-1 agents); this is because our assumption that an agent's utility only depends on the *aggregate* decisions by groups other than the one they belong to implies certain properties of the dependence structure. For the remainder of this paper we will proceed with a given multi-scale structure defined above, while in Appendix G we outline a set of conditions on a graphical game G that allows us to represent it in a (non-trivial) multi-scale fashion.

Since the resulting multi-scale network is strictly hierarchical, we can define a *direct supervisor* of agent $a_i^{(l)}$ in level- l to be the agent $a_k^{(l+1)}$ corresponding to the level- $(l + 1)$ group k that the former belongs to. Similarly, two agents who belong in the same level- l group k are (level- l) *group mates*. Finally, note that any level-1 agent a_i belongs to exactly one group in each level l . We index a level- l group to which a_i belongs by k_{il} .

In order to capture the agent dependence on aggregate actions, we define an *aggregation function* $\sigma_k^{(l)}$ for each level- l group k that maps individual actions of group members to \mathbb{R} (a *group strategy*). Specifically, consider a level- l group $S_k^{(l)}$ with level- $(l - 1)$ agents in this group playing a strategy profile $\mathbf{x}_{S_k^{(l)}}$. The (scalar) group strategy, which is also the strategy for the corresponding level- $(l + 1)$ agent, is determined by the aggregation function,

$$x_k^{(l)} = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}). \quad (5)$$

A natural example of this is linear (e.g., agents respond to total levels of violence by other criminal organizations): $\sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}) = \sum_{i \in S_k^{(l)}} x_i^{(l)}$.

The L -level structure above is captured strategically by introducing structure into the utility functions of agents. Let $I_{k_{il}}$ denote the set of neighbors of level- l group k to which level-1 agent a_i belongs; i.e., this is the set of level- l groups that interact with agent a_i 's group. This level-1 agent's utility function can be decomposed as follows:

$$u_i(x_i, \mathbf{x}_{-i}) = \sum_{l=1}^L u_{k_{il}}^{(l)}(x_{k_{il}}^{(l)}, \mathbf{x}_{I_{k_{il}}}^{(l)}). \quad (6)$$

In this definition, the level- l strategies $x_k^{(l)}$ are implicitly functions of the level-1 strategies of agents that comprise the group, per a recursive application of Eqn. (5). Consequently, the utility is an additive function of the hierarchy of group-level components for increasingly (with l) abstract group of agents. Note that conventional network games are a special case with only a single level ($L = 1$).

To illustrate, if we consider just two levels (a collection of individuals and groups to which they directly belong), the utility function of each agent a_i is a sum of two components:

$$u_i(x_i, \mathbf{x}_{-i}) = u_{k_{i1}}^{(1)}\left(x_{k_{i1}}^{(1)}, \mathbf{x}_{I_{k_{i1}}}^{(1)}\right) + u_{k_{i2}}^{(2)}\left(x_{k_{i2}}^{(2)}, \mathbf{x}_{I_{k_{i2}}}^{(2)}\right).$$

In the first component, $x_{k_{i1}}^{(1)} = x_i$, since level-1 groups correspond to individual agents, whereas $\mathbf{x}_{I_{k_{i1}}}^{(1)}$ is the strategy profile of i 's neighbors *belonging to the same group as i* , given by $E^{(1)}$. The second utility component now depends only on the aggregate strategy $x_{k_{i2}}^{(2)}$ of the group to which i belongs, as well as the aggregate strategies of the groups with which i 's group interacts, given by $E^{(2)}$.

4 Algorithms and Analysis

Consider the BRD algorithm (formalized in Algorithm 1) in which we iteratively select an agent who plays a best response to the strategy of the rest from the previous iteration.

ALGORITHM 1: BRD Algorithm

```

Initialize the game,  $t = 0, x_i(0) = (\mathbf{x}_0)_i, i = 1, \dots, N;$ 
while not converged do
  for  $i = 1:N$  do
     $x_i(t+1) = BR_i(\mathbf{x}_{-i}(t), u_i)$ 
  end
   $t \leftarrow t + 1$ 
end

```

The conventional BRD algorithm operates on the “flattened” utility function which evaluates utilities explicitly as functions of the strategies played by all agents $a_i \in I$. Our goal henceforth is to develop algorithms that take advantage of the special multi-scale structure and enable significantly better scalability than standard BRD, while preserving the convergence properties of BRD.

4.1 Taking Advantage of Multi-Scale Utility Representation

The simplest way to take advantage of the multi-scale representation is to directly leverage the structure of the utility function in computing best responses. Specifically, the multi-scale utility function is more compact than one that explicitly accounts for the strategies of all neighbors of i (which includes *all* of the players in groups other than the one i belongs to). This typically results in a direct computational benefit to computing a best response. For example, in a game with a linear best response, this can result in an exponential reduction in the number of linear operations.

The resulting algorithm, *Multi-Scale Best-Response Dynamics (MS-BRD)*, which takes advantage of our utility representation is formalized as Algorithm 2. The main difference from BRD is that it explicitly uses the multi-scale utility representation: in each iteration, it updates the aggregated strategies at all levels for the groups to which the most recent best-responding agent belongs. Since MS-BRD simply

ALGORITHM 2: Multi-Scale BRD (MS-BRD)

```

Initialize the game,  $t = 0, x_i^{(1)}(0) = (\mathbf{x}_0)_i, i = 1, \dots, N$ 
for  $l = 2:L$  do
  for  $k = 1:N^{(l)}$  do
     $\mathbf{x}_k^{(l)}(0) = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}(0));$ 
  end
end
while not converged do
  for  $i = 1:N$  (Level-1) do
     $x_i^{(1)}(t+1) = BR_i(\mathbf{x}_{-i}^{(1)}(t), u_i)$ 
  end
  for  $l = 2:L$  do
    for  $k = 1:N^{(l)}$  do
       $\mathbf{x}_k^{(l)}(t+1) = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}(t+1));$ 
    end
  end
   $t \leftarrow t + 1;$ 
end

```

performs operations identical to BRD but efficiently, its convergence is guaranteed under the same conditions (see Theorem 1). Next, we present iterative algorithms for computing NE that take further advantage of the multi-scale structure, and study their convergence.

4.2 Taking Advantage of Multi-Scale Strategic Dependence Structure

In order to take full advantage of the multi-scale game structure, we now aim to develop algorithms that treat groups explicitly as agents, with the idea that iterative interactions among these can significantly speed up convergence. Of course, in our model groups are not actual agents in the game: utility functions are only defined for agents in level 1. However, note that we already have well-defined group strategies – these are just the aggregations of agent strategies at the level immediately below, per the aggregation function (5). Moreover, we have natural utilities for groups as well: we can use the corresponding group-level component of the utility of any agent in the group (note that these are identical for all group members in Eqn. (6)). However, using these as group utilities will in fact not work: since ultimately the game is only among the agents in level 1, equilibria of all of the games at more abstract levels *must be consistent with equilibrium strategies in level 1*. On the other hand, we need to enforce consistency only between neighboring levels, since that fully captures the across-level interdependence induced by the aggregation function. Therefore, we define the following *pseudo-utility functions* for agents at levels other than 1, with agent k in level l corresponding to a subset of agents from level $l - 1$:

$$\hat{u}_k^{(l)} = u_k^{(l)}\left(x_k^{(l)}, \mathbf{x}_{I_k}^{(l)}\right) - L_k^{(l,l-1)}\left(x_k^{(l)}, \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}})\right) - L_k^{(l,l+1)}\left(\sigma_k^{(l+1)}(\mathbf{x}_{S_k^{(l+1)}}), x_k^{(l+1)}\right). \quad (7)$$

The first term is the level- l component of the utility of any level-1 agent in group k . The second and third terms model the inter-level inconsistency loss that penalizes a level- l agent $a_k^{(l)}$, where $L_k^{(l,l+1)}$ and $L_i^{(l,l-1)}$ penalize its inconsistency with the level- $(l+1)$ and level- $(l-1)$ entities respectively. In general, $L_k^{(l,l+1)}$ is a different function from $L_k^{(l+1,l)}$; we elaborate on this further below.

The central idea behind the second algorithm we propose is simple: in addition to iterating best response steps at level 1, we now interleave them with best response steps taken by agents at higher levels, which we can since strategies and utilities of these pseudo-agents are well defined. This algorithm is similar to the augmented Lagrangian method in optimization theory, where penalty terms are added to relax an equality constraint and turn the problem into one with separable operators. We can decompose this type of problem into smaller subproblems and solve the subproblems sequentially using the alternating direction method (ADM) (Yuan and Li, 2011; Bnouhachem, Benazza, and Khalfaoui, 2013). The games at adjacent levels are coupled through the equality constraints on their action profiles given by Eqn (5), and the penalty functions are updated before starting a new iteration. The full algorithm, which we call *Separated Hierarchical BRD (SH-BRD)*, is provided in Algorithm (3).

The penalty updating rule in iteration t of Algorithm (3) is:

1. For $l = 2, \dots, L, i = 1, \dots, N^{(l)}$

$$\begin{aligned} & L_i^{(l,l-1)} \left(x_i^{(l)}, \sigma_i^{(l)}(\mathbf{x}_{S_i^{(l)}}(t+1)) \right) \\ &= h_i^{(l)} \left[x_i^{(l)} - \sigma_i^{(l)}(\mathbf{x}_{S_i^{(l)}}(t+1)) + \lambda_i^{(l)}(t) \right]^2. \end{aligned} \quad (8)$$

2. For $l = 1, \dots, L-1; i = 1, \dots, N^{(l)}$, where $a_i^{(l)} \in S_k^{(l+1)}$

$$\begin{aligned} & L_k^{(l,l+1)} \left(\sigma_k^{(l+1)}(\mathbf{x}_{S_k^{(l+1)}}), x_k^{(l+1)}(t) \right) \\ &= h_k^{(l+1)} \left[\sigma_k^{(l+1)}(\mathbf{x}_{S_k^{(l+1)}}) - x_k^{(l+1)}(t) - \lambda_k^{(l+1)}(t) \right]^2. \end{aligned} \quad (9)$$

3. For $l = 2, \dots, L, i = 1, \dots, N^{(l)}$

$$\begin{aligned} & \lambda_i^{(l)}(t+1) \\ &= \lambda_i^{(l)}(t) - h_i^{(l)} \left[\sigma_i^{(l)}(\mathbf{x}_{S_i^{(l)}}(t+1)) - x_i^{(l)}(t+1) \right]. \end{aligned} \quad (10)$$

When updating, all other variables are treated as fixed, and $\lambda^{(l)}(0), h_i^{(l)} > 0$ are chosen arbitrarily.

Unlike MS-BRD, the convergence of the SH-BRD algorithm is non-trivial. To prove it, we exploit a connection between this algorithm and Structured Variational Inequalities (SVI) with separable operators (He, 2009; He and Yuan, 2012; Bnouhachem, Benazza, and Khalfaoui, 2013). To formally state the convergence result, we need to make several explicit assumptions.

ALGORITHM 3: Separated Hierarchical BRD (SH-BRD)

Initialize the game, $t = 0, x_i^{(1)}(0) = (\mathbf{x}_0)_i, i = 1, \dots, N^{(0)}$

```

for  $l = 2:L$  do
  for  $k = 1:N^{(l)}$  do
     $\mathbf{x}_k^{(l)}(0) = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}(0));$ 
  end
end
while not converged do
  for  $l = 1:L$  do
    for  $i = 1:N^{(l)}$  ( $l$  to  $l-1$  Penalty Update, if  $l > 1$ ) do
      Update  $L_i^{(l,l-1)}$ 
    end
    for  $i = 1:N^{(l)}$  ( $l$  to  $l+1$  Penalty Update, if  $l < L$ ) do
      Update  $L_k^{(l,l+1)}$ , where  $a_i^{(l)} \in S_k^{(l+1)}$ 
    end
    for  $i = 1:N^{(l)}$  (Best Response) do
       $x_i^{(l)}(t+1) = BR_i \left( \sigma_i^{(l)}(\mathbf{x}_{S_i^{(l)}}(t+1)), \right.$ 
       $\left. \mathbf{x}_{I_i}^{(l)}(t), x_k^{(l+1)}(t), \hat{u}_i^{(l)} \right)$ 
    end
  end
   $t \leftarrow t + 1;$ 
end

```

Assumption 1. The functions $u_i^{(l)}, \forall l = 1, \dots, L, \forall i = 1, \dots, N^{(l-1)}$ are twice continuously differentiable.

Assumption 2. $-\nabla_{x_i^{(l)}} u_i^{(l)}$ are monotone $\forall l = 1, \dots, L, \forall i = 1, \dots, N^{(l-1)}$. The solution set of $\nabla_{x_i^{(l)}} u_i^{(l)} = 0, \forall l = 1, \dots, L, \forall i = 1, \dots, N^{(l-1)}$ is nonempty, with solutions in the interior of the action spaces.

Let $F^{(l)}$ be defined as in Equation (3) for each level- l pseudo-utility.

Assumption 3. $F^{(l)}$ satisfy the P_Υ condition.

Note that these assumptions directly generalize the conditions required for the convergence of BRD to our multi-scale pseudo-utilities. The following theorem formally states that SH-BRD converges to a NE for 2-level games.

Theorem 2. Suppose $L = 2$. If Assumptions 1 and 3 hold, SH-BRD converges to a NE, which is unique.

The full proof of this theorem, which makes use of the connection between SH-BRD and SVI, is provided in the Supplement due to space constraint. The central issue, however, is that there are no established convergence guarantees for ADM-based algorithms for SVI with 3 or more separable operators. Alternative algorithms for SVI can extend to the case of 3 operators using parallel operator updates with regularization terms, but no approaches exist that can handle more than 3 operators (He, 2009). We thus propose an algorithm for iteratively solving multi-scale games that uses the general idea from SH-BRD, but packs all levels into two meta-levels. The two meta-levels each has to be comprised

of consecutive levels. For example, if we have 5 levels, we can have $\{1, 2, 3\}$ and $\{4, 5\}$ combinations, but not $\{1, 2, 4\}$ and $\{3, 5\}$. Upon grouping levels together to obtain a meta-game with only two meta-levels, we can apply what amounts to a 2-level version of the SH-BRD. This yields an algorithm, which we call *Hybrid Hierarchical BRD (HH-BRD)*, that now provably converges to a NE for an arbitrary number of levels L given assumptions 1-3.

As presenting the general version of HH-BRD involves cumbersome notation, we illustrate the idea by presenting it for a 4-level game (Algorithm 4). The fully general version is deferred to the Supplement. In this example, the objectives of the meta-levels are defined as

$$\begin{aligned}\hat{u}_i^{(sl_1)} &= u_i^{(1)} + u_{k_{i2}}^{(2)} - L_{k_{i3}}^{(sl_1, sl_2)} \left(\sigma_{k_{i3}}^{(3)}(\mathbf{x}_{S_{k_{i3}}^{(3)}}), x_{k_{i3}}^{(3)} \right), \\ \hat{u}_{k_{i3}}^{(sl_2)} &= u_{k_{i3}}^{(3)} + u_{k_{i4}}^{(4)} - L_{k_{i3}}^{(sl_2, sl_1)} \left(x_{k_{i3}}^{(3)}, \sigma_{k_{i3}}^{(3)}(\mathbf{x}_{S_{k_{i3}}^{(3)}}) \right).\end{aligned}$$

ALGORITHM 4: Hybrid Hierarchical BRD

Initialize the game, $t = 0, x_i^{(1)}(0) = (\mathbf{x}_0)_i, i = 1, \dots, N^{(0)}$

```

for  $l = 2:4$  do
  for  $k = 1:N^{(l)}$  do
     $\mathbf{x}_k^{(l)}(0) = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}(0));$ 
  end
end
while not converged do
  for  $k = 1:N^{(3)}$  (Meta-Level-1 Penalty Update) do
    Update  $L_k^{(sl_1, sl_2)}$ 
  end
  for  $i = 1 : N^{(1)}$  (Level-1) do
     $x_i^{(1)}(t+1) =$ 
     $BR_i \left( \mathbf{x}_{I_i}^{(1)}(t), \mathbf{x}_{I_{k_{i2}}}^{(2)}(t), \mathbf{x}_{k_{i3}}^{(3)}(t), \hat{u}_i^{(sl_1)} \right)$ 
  end
  for  $j = 1:N^{(2)}$  (Level-2) do
     $\mathbf{x}_j^{(2)}(t+1) = \sigma_j^{(2)}(\mathbf{x}_{S_j^{(2)}}(t+1))$ 
  end
  for  $k = 1:N^{(3)}$  (Meta-Level-2 Penalty Update) do
    Update  $L_k^{(sl_2, sl_1)}$ 
  end
  for  $k = 1 : N^{(3)}$  (Level-3) do
     $x_k^{(3)}(t+1) = BR_k \left( \sigma_k^{(3)}(\mathbf{x}_{S_k^{(3)}}(t+1)), \mathbf{x}_{I_k}^{(3)}(t), \right.$ 
     $\left. \mathbf{x}_{-p}^{(4)}(t), \hat{u}_k^{(sl_2)} \right), (a_k^{(3)} \in S_p^{(4)})$ 
  end
  for  $p = 1:N^{(4)}$  (Level-4) do
     $\mathbf{x}_p^{(4)}(t+1) = \sigma_p^{(4)}(\mathbf{x}_{S_p^{(4)}}(t+1))$ 
  end
   $t \leftarrow t + 1;$ 
end

```

Theorem 3. *Suppose Assumptions 1-3 hold Then HH-BRD finds the unique NE.*

Proof Sketch. We first “flatten” the game within each meta-level to obtain an effective 2-level game. We then use Theorem 2 to show this 2-level game converges to the unique NE of the game under SH-BRD. Finally, we prove that SH-BRD and HH-BRD have the same trajectory given the same initialization, thus establishing the convergence for HH-BRD. For full proof see Supplement, Appendix D. \square

HH-BRD combines the advantages of both MS-BRD and SH-BRD: not only does it exploit the sparsity embedded in the network topology, but it also avoids the convergence problem of SH-BRD when the number of levels is higher than three. Indeed, there is a known challenge in the related work on structured variational inequalities that convergence is difficult when we involve three or more operators (He, 2009), which we leverage for our convergence results, with operators mapping to levels in our multi-scale game representation. One may be concerned that HH-BRD pseudocode appears to involve greater complexity (and more steps) than SH-BRD. However, this does not imply greater algorithmic complexity, but is rather due to our greater elaboration of the steps within each super level. Indeed, as our experiments below demonstrate, the superior theoretical convergence of HH-BRD also translates into a concrete computational advantage of this algorithm.

5 Numerical Results and Analysis

In this section, we numerically compare the three algorithms introduced in Section 4, as well as the conventional BRD. We only consider settings which satisfy Assumptions 1-3; consequently, we focus comparison on computational costs. We use two measures of computational cost: floating-point operations (FLOPs) in the case of games with a linear best response (a typical measure for such settings), and CPU time for the rest. All experiments were performed on a machine with A 6-core 2.60/4.50 GHz CPU with hyperthreaded cores, 12MB Cache, and 16GB RAM.

Games with a Linear Best Response (GLBRs) GLBRs (Bramoullé, Kranton, and D’amours, 2014; Candogan, Bimpikis, and Ozdaglar, 2012; Miura-Ko et al., 2008) feature utility functions such that an agent’s best response is a linear function of its neighbors’ actions. This includes quadratic utilities of the form

$$u_i(x_i, x_{I_i}) = a_i + b_i x_i + \left(\sum_{j \in I_i} g_{ij} x_j \right) x_i - c_i x_i^2, \quad (11)$$

since an agent’s best response is:

$$BR_i(x_{I_i}, u_i) = \frac{\sum_{j \in I_i} g_{ij} x_j}{2c_i} - b_i.$$

We consider a 2-level GLBR and compare three algorithms: BRD (baseline), MS-BRD, and HS-BRD (note that in 2-level games, HH-BRD is identical to HS-BRD, and we thus don’t include it here). We construct random 2-level games with 10,000 agents and 100 groups (level-2 agents), with utility functions based on Equation (11). Specifically,

we generalize this utility so that Equation (11) represents only the level-1 portion, $u_i^{(1)}$, and let the level-2 utilities be

$$u_k^{(2)}(x_k, \mathbf{x}_{I_k}) = x_k^{(2)} \sum_{p \neq k} v_{kp} x_p^{(2)}$$

for each group k . At every level, the existence of a link between two agents follows the Bernoulli distribution where $P_{exist} = 0.1$. If a link exists, we then generate a parameter for it. The parameters of the utility functions are sampled uniformly in $[0, 1]$ without requiring symmetry. Please refer to Appendix E and E.1 for further details. Results comparing BRD, MS-BRD, and SH-BRD are shown in Table 5. We observe dramatic improvement in the scalability of using MS-BRD compared to conventional BRD. This improvement stems from the representational advantage provided by multi-scale games compared to conventional graphical games (since without the multi-scale representation, we have to use the standard version of BRD for equilibrium computation). We see further improvement going from MS-BRD to SH-BRD which makes algorithmic use of the multi-scale representation.

Size	BRD	MS-BRD	SH-BRD
30^2	$(2.51 \pm 0.18) \times 10^6$	$(1.03 \pm 0.07) \times 10^5$	$(9.81 \pm 0.81) \times 10^4$
50^2	$(2.53 \pm 0.18) \times 10^7$	$(5.33 \pm 0.04) \times 10^5$	$(4.35 \pm 0.07) \times 10^5$
100^2	$(4.46 \pm 0.32) \times 10^8$	$(4.36 \pm 0.31) \times 10^6$	$(3.56 \pm 0.29) \times 10^6$
200^2	$(6.73 \pm 0.58) \times 10^9$	$(3.48 \pm 0.29) \times 10^7$	$(2.79 \pm 0.21) \times 10^7$
500^2	$(2.84 \pm 0.21) \times 10^{11}$	$(5.69 \pm 0.41) \times 10^8$	$(4.04 \pm 0.29) \times 10^8$

Table 1: Convergence and complexity (flops) comparison with linear best response under multiple initialization.

Games with a Non-Linear Best Response Next, we study the performance of the proposed algorithms in 2- and 3-level games, with the same number of groups in each level (we systematically vary the number of groups). Since HS-BRD and HH-BRD are identical in 2-level games, the latter is only used in 3-level games. All results are averaged over 30 generated sample games. The non-linear best response fits a much broader class of utility functions than the linear best response. The best responses generally don't have closed-form representations. In this case, we can't use linear equations to find the best response and instead have to apply gradient-based methods. In our instances, the utility with non-linear best responses is generated by adding an exponential cost term to the utility function used in GLBRs. Please refer to Appendix E and E.2 for further details.

Table 2 shows the CPU time comparison between all algorithms. The scalability improvements from our proposed algorithms are substantial, with orders of magnitude speedup in some cases (e.g., from ~ 25 minutes for the BRD baseline, down to ~ 12 seconds for SH-BRD for games with 10K agents). Furthermore, BRD fails to solve instances with 250K agents, which can be solved by SH-BRD in ~ 42 min. Again, we separate here the representational advantage of

Size	BRD	MS-BRD	SH-BRD
30^2	1.50 ± 0.05	1.02 ± 0.02	0.54 ± 0.01
50^2	26.70 ± 0.36	3.70 ± 0.14	1.81 ± 0.04
100^2	1512 ± 9 s	23.81 ± 0.69	12.10 ± 0.13
200^2	> 18000	287.2 ± 5.4	133.6 ± 2.5
500^2	nan	5485 ± 13 s	2524 ± 10

Table 2: CPU times on a single machine on 2-Level games with general best response functions; all times are in seconds.

multi-scale games, illustrated by MS-BRD, and algorithmic advantage that comes from SH-BRD. Note that SH-BRD, which takes full advantage of the multi-scale structure, also exhibits significant improvement over MS-BRD, yielding a factor of 2-3 reduction in runtime.

Size	BRD	MS-BRD	SH-BRD
30^2	1.21 ± 0.04	0.63 ± 0.01	0.037 ± 0.003
50^2	23.88 ± 0.16	1.99 ± 0.04	0.079 ± 0.004
100^2	1461 ± 14	15.49 ± 0.24	0.304 ± 0.006
200^2	> 18000	192.0 ± 1.2	1.87 ± 0.05
500^2	nan	4258 ± 56 s	28.79 ± 0.37

Table 3: CPU times on a single machine for 2-Level, linear/nonlinear best-response games; all times are in seconds.

Our next set of experiments involves games in which level-1 utility has a linear best response, but level-2 utility has a non-linear best response. The results are shown in Table 3. We see an even bigger advantage of SH-BRD over the others: it is now typically orders of magnitude faster than even MS-BRD, which is itself an order of magnitude faster than BRD. For example, in games with 250K agents, in which BRD fails to return a solution, MS-BRD takes more than 1 hour to find a solution, whereas SH-BRD finds a solution in under 30 seconds.

Size	BRD	MS-BRD	SH-BRD	HH-BRD
10^3	1.23 ± 0.03	0.59 ± 0.01	0.76 ± 0.03	0.43 ± 0.02
20^3	696.0 ± 8.7	3.78 ± 0.09	6.05 ± 0.08	3.35 ± 0.09
30^3	> 18000	15.70 ± 0.11	25.13 ± 0.14	13.39 ± 0.11
50^3	nan	68.59 ± 0.75	138.8 ± 1.1	57.98 ± 0.69
100^3	nan	1126 ± 6	2343 ± 21	877.1 ± 11.5

Table 4: CPU times in seconds on a single machine on 3-Level, general best response games; all times are in seconds.

Finally, Table 4 presents the results of HH-BRD in games with > 2 levels compared to SH-BRD, which does not provably converge in such games. In this case, HH-BRD outperforms the other alternatives, with up to 22% improvement over MS-BRD; indeed, we find that SH-BRD is considerably worse even than MS-BRD.

6 Conclusions and Future Directions

We proposed a novel representation of games that have a multi-scale network structure. These generalize network games, but with special structures that agent utilities are additive across the levels of hierarchy, with utility at each level depending only on the aggregate strategies of other groups. We present several iterative algorithms that make use of the multi-scale game structure, and show that they converge to a pure strategy Nash equilibrium under similar conditions as for best response dynamics in network games. Our experiments demonstrate that the proposed algorithms can yield orders of magnitude scalability improvement over conventional best response dynamics. Our multi-scale algorithms can reveal to what extent one's group affiliation impacts one's strategic decision making, and how strategic interactions among groups impact strategic interactions among individuals.

While the issue of multi-scale networks abounds in the network science literature (e.g., hierarchical clustering, etc.), the "multi-scale" part is primarily concerned with community structure in networks, rather than modeling how *how communities interact*, which is critical for us in describing a formal multi-scale structure for games. Thus a very important future direction is to identify and obtain relevant field data for experiments, and create realistic benchmarks for multi-scale games. This would involve identifying ways to obtain data about how communities (and not just individuals) interact. Once we have the ability to collect data about interactions at multiple scales (e.g., among members and among groups), we can apply our algorithms to such multi-scale networks. To use criminal networks (criminal organizations and their members) as an example, given game models constructed with the help of domain expertise, we can:

1. compute equilibria predicting, say, criminal activity as a function of structural changes to organizations;
2. infer utility models from observational data at multiple scales;
3. study policies (including strengthening or weakening connections between agents or groups, endowing agents/groups with more resources (lower costs of effort), etc.) that would induce more desirable equilibrium outcomes.

Acknowledgment

This work is supported by the NSF under grants CNS-1939006, CNS-2012001, IIS-1905558 (CAREER) and by the ARO under contract W911NF1810208.

References

- Acemoglu, D.; Carvalho, V. M.; Ozdaglar, A.; and Tahbaz-Salehi, A. 2012. The network origins of aggregate fluctuations. *Econometrica* 80(5):1977–2016.
- Allouch, N. 2015. On the private provision of public goods on networks. *Journal of Economic Theory* 157:527–552.
- Alvari, H.; Hajibagheri, A.; and Sukthankar, G. 2014. Community detection in dynamic social networks: A game-theoretic approach.
- Bnouchem, A.; Benazza, H.; and Khalfaoui, M. 2013. An inexact alternating direction method for solving a class of structured variational inequalities. *Applied Mathematics and Computation* 219:7837–7846.
- Bramoullé, Y.; Kranton, R.; and D'amours, M. 2014. Strategic interaction and networks. *American Economic Review* 104(3):898–930.
- Bramoullé, Y., and Kranton, R. 2016. *Games Played on Networks*.
- Buckley, E., and Croson, R. 2006. Income and wealth heterogeneity in the voluntary provision of linear public goods. *Journal of Public Economics* 90(4-5):935–955.
- Candogan, O.; Bimpikis, K.; and Ozdaglar, A. 2012. Optimal pricing in networks with externalities. *Operations Research* 60(4):883–905.
- Clauset, A.; Moor, C.; and Newman, M. 2008. Hierarchical structure and the prediction of missing links in networks. *Nature* 453:98–101.
- Gabay, D., and Mercier, B. 1976. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications* 2:17–40.
- Girvan, M., and Newman, M. E. J. 2002. Community structure in social and biological networks. *Proc. Natl. Acad. Sci.* 99:7821–7826.
- Glowinski, R., and Oden, J. 1985. Numerical methods for nonlinear variational problems. *Journal of Applied Mechanics* 52:739–.
- He, B.-S., and Yuan, X. 2012. On the $o(1/t)$ convergence rate of alternating direction method.
- He, B.-S.; Yang, H.; and Wang, S. 2000. Alternating direction method with self-adaptive penalty parameters for monotone variational inequalities. *Journal of Optimization Theory and Applications* 106:337–356.
- He, B.-S. 2009. Parallel splitting augmented lagrangian methods for monotone structured variational inequalities. *Computational Optimization and Applications* 42:195–212.
- Hota, A. R., and Sundaram, S. 2018. Interdependent security games on networks under behavioral probability weighting. *IEEE Transactions on Control of Network Systems* 5(1):262–273.
- Jackson, M. O., and Zenou, Y. 2015. Games on networks. In *Handbook of game theory with economic applications*, volume 4. Elsevier. 95–163.

- Kearns, M. J.; Littman, M. L.; and Singh, S. P. 2001. Graphical models for game theory. In *Conference in Uncertainty in Artificial Intelligence*, 253–260.
- Khalili, M. M.; Zhang, X.; and Liu, M. 2019. Public good provision games on networks with resource pooling. In *Network Games, Control, and Optimization*. Springer. 271–287.
- La, R. J. 2016. Interdependent security with strategic agents and cascades of infection. *IEEE/ACM Transactions on Networking* 24(3):1378–1391.
- Lions, P., and Mercier, B. 1979. Mercier, b.: Splitting algorithms for the sum of two nonlinear operators. *siam j. numer. anal.* 16(6), 964–979. *Siam Journal on Numerical Analysis - SIAM J NUMER ANAL* 16:964–979.
- Mcsweeney, P.; Mehrotra, K.; and Oh, J. 2017. *Game-Theoretic Framework for Community Detection*. 1–16.
- Miura-Ko, R. A.; Yolken, B.; Mitchell, J.; and Bambos, N. 2008. Security decision-making among interdependent organizations. In *2008 21st IEEE Computer Security Foundations Symposium*, 66–80.
- Naghizadeh, P., and Liu, M. 2017. On the uniqueness and stability of equilibria of network games. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 280–286. IEEE.
- Newman, M. 2004. Detecting community structure in networks. *Eur Phys J* 38.
- Parise, F., and Ozdaglar, A. 2019. A variational inequality framework for network games: Existence, uniqueness, convergence and sensitivity analysis. *Games and Economic Behavior*.
- Scutari, G.; Facchinei, F.; Pang, J.-S.; and Palomar, D. P. 2014. Real and complex monotone communication games. *IEEE Transactions on Information Theory* 60(7):4197–4231.
- Tseng, P. 1990. Applications of a splitting algorithm to decomposition in convex programming and variational inequalities. *Mathematical Programming* 48:249–263.
- Vorobeychik, Y., and Letchford, J. 2015. Securing interdependent assets. *Journal of Autonomous Agents and Multiagent Systems* 29:305–333.
- Yu, S.; Zhou, K.; Brantingham, P. J.; and Vorobeychik, Y. 2020. Computing equilibria in binary networked public goods games. In *AAAI Conference on Artificial Intelligence*, 2310–2317.
- Yuan, X.-M., and Li, M. 2011. An lqp-based decomposition method for solving a class of variational inequalities. *SIAM Journal on Optimization* 21:1309–1318.

Appendices

A Structured Variational Inequalities

A structured variational inequality SVI_n arises when a VI problem has n separable operators. This is used to analyze our game under the multi-scale perspective described in Section 3.

We now introduce a particular type of SVI_2 relevant to our model. Suppose the N level-1 agents form M disjoint groups in the game and S_j denotes the j th level-1 group, whereby $i \in S_j$ denotes that a_i is a member of S_j . Consider the following utility function of a_i :

$$u_i(x_i, \mathbf{x}_{-i}, y_j, \mathbf{y}_{-j}) = u_i^{(1)}(x_i, \mathbf{x}_{-i}) + u_j^{(2)}(y_j, \mathbf{y}_{-j}), \quad (12)$$

where $\mathbf{x} \in \mathbf{R}^N$ denotes the level-1 action profile and $\mathbf{y} \in \mathbf{R}^M$ denotes the level-2 action profile, and $A\mathbf{x} + \mathbf{y} = \mathbf{0}$, for

$$A_{ji} = \begin{cases} -1, & \text{if } i \in S_j \\ 0, & \text{else} \end{cases}, j = 1, \dots, M, i = 1, \dots, N.$$

Thus $A\mathbf{x} + \mathbf{y} = \mathbf{0}$ is equivalent to $y_j = \sum_{i \in S_j} x_i$. We say \mathbf{x} and \mathbf{y} are two separated operators, and define

$$\begin{aligned} F^{(1)}(\mathbf{x}) &:= \left(-\nabla_{x_i} u_i^{(1)}(\mathbf{x}) \right)_{i=1}^N, \quad x_i \in K_i^{(1)}, \\ F^{(2)}(\mathbf{y}) &:= \left(-\nabla_{y_j} u_j^{(2)}(\mathbf{y}) \right)_{j=1}^M, \quad y_j \in K_j^{(2)}, \\ K^{(1)} &= \prod_{i=1}^N K_i^{(1)}, K^{(2)} = \prod_{j=1}^M K_j^{(2)}, \bar{K} = K^{(1)} \times K^{(2)}, \\ \mathbf{v} = \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} &\in \bar{K}, \quad \bar{F}(\mathbf{v}) = \begin{bmatrix} F^{(1)}(\mathbf{x}) \\ F^{(2)}(\mathbf{y}) \end{bmatrix}. \end{aligned} \quad (13)$$

Define $\Omega = \{v \in \bar{K} | A\mathbf{x} + \mathbf{y} = \mathbf{0}\}$. Then the $\text{VI}(\Omega, \bar{F})$ problem is to find $v^* \in \Omega$, such that: $(v - v^*)^T \bar{F}(v) \geq 0, \forall v \in \Omega$. This problem is equivalent to the SVI_2 problem $\text{VI}(\mathcal{W}, Q)$ defined in Eqn (14)

$$(\omega - \omega^*)^T Q(\omega) \geq 0, \forall \omega \in \mathcal{W}, \quad (14)$$

where, $\mathcal{W} = \bar{K} \times \mathbf{R}^M$ and

$$\omega = \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \\ \lambda \end{pmatrix}, Q(\omega) = \begin{pmatrix} F^{(1)}(\mathbf{x}) - A^T \lambda \\ F^{(2)}(\mathbf{y}) - \lambda \\ A\mathbf{x} + \mathbf{y} \end{pmatrix}. \quad (15)$$

It is easy to see that if we use $\sum_{i \in S_j} x_i$ to replace y_j , then we again have a single operator \mathbf{x} and can construct a $\text{VI}(K, F)$ as outlined in Section ???. There is a one-to-one mapping between a solution \mathbf{x}^* to $\text{VI}(K, F)$ and a solution $\omega^* = (\mathbf{x}^*, -A\mathbf{x}^*, \lambda^*)$ to $\text{VI}(\mathcal{W}, Q)$. Therefore, solving either $\text{VI}(K, F)$ or $\text{VI}(\mathcal{W}, Q)$ finds the set of NEs.

B Uniqueness of NE

We will introduce some special matrices before we move on to the sufficient conditions for the uniqueness of NE.

Definition 2. *Some special matrices:*

1. *P-matrix:* A square matrix is a P-matrix if all its principal components have positive determinant
2. *Z-matrix:* A square matrix is a Z-matrix if all its off-diagonal components are nonpositive
3. *M-matrix:* An M-matrix is a Z-matrix whose eigenvalues' real parts are nonnegative
4. *L-matrix:* An L-matrix is a Z-matrix whose diagonal elements are nonnegative

For an arbitrary mapping $F : \mathbf{R}^N \rightarrow \mathbf{R}^N$, we denote the Jacobian of $F(\mathbf{x})$ as $JF(\mathbf{x})$. And then $\nabla_j F_i = [JF(\mathbf{x})]_{ij}$

Checking if a matrix is P-matrix or not is still not trivial, and we can look at the spectral radius of a matrix instead.

Theorem 4. The P_Γ condition:

We define the Γ matrix generated from F as follows

$$\Gamma(F) = \begin{bmatrix} 0 & -\frac{\beta_{1,2}(F)}{\alpha_1(F)} & \dots & -\frac{\beta_{1,N}(F)}{\alpha_1(F)} \\ -\frac{\beta_{2,1}(F)}{\alpha_2(F)} & 0 & \dots & -\frac{\beta_{2,N}(F)}{\alpha_2(F)} \\ \vdots & \vdots & \ddots & \vdots \\ -\frac{\beta_{N,1}(F)}{\alpha_N(F)} & -\frac{\beta_{N,2}(F)}{\alpha_N(F)} & \dots & 0 \end{bmatrix}, \quad (16)$$

if the spectral radius $\rho(\Gamma(F)) = \|\Gamma(F)\|_2 < 1$, then we say F satisfies the P_Γ condition. Then P_Γ condition $\Leftrightarrow P_\Upsilon$ condition and $\text{VI}(K, F)$ has a unique solution.

In Scutari et al. (2014), the authors mentioned that the P_Υ captures ‘‘some kind of diagonal dominance’’. In fact, the strong diagonal dominance(s.d.d) or weakly chained diagonal dominance(w.c.d.d) of Υ can be an easier yet sufficient condition to check.

Theorem 5. *If Υ is s.d.d or w.c.d.d, the NE is unique, since*

- Υ is an s.d.d L-matrix
- $\Rightarrow \Upsilon$ is a w.c.d.d L-matrix
- $\Leftrightarrow \Upsilon$ is a nonsingular weakly diagonally dominant(w.d.d) L-matrix
- $\Leftrightarrow \Upsilon$ is a nonsingular w.d.d M-matrix
- $\Rightarrow \Upsilon$ is a P-matrix

Also, when Υ is s.d.d, Γ is a (right, row) substochastic matrix and thus $\rho(\Gamma) < 1$ trivially holds and the NE is unique.

The P_Υ condition guarantees both the uniqueness of NE and the convergence of BRD. Please refer to Parise and Ozdaglar (2019) for more conditions on the uniqueness.

C Proof of Theorem 2

Proof. This algorithm is designed to solve the SVI problem presented in Eqn (14) and (15). We denote $H = \frac{1}{2} \mathbf{diag}(\mathbf{h})$, and the norm $\|\mathbf{x}\|_G$, where $G \succ \mathbf{0}$ as

$$\|\mathbf{x}\|_G = \mathbf{x}^T G \mathbf{x}.$$

For simplicity reason, we will use \mathbf{x} and \mathbf{y} to replace $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ in the remainder of the proof.

We can rewrite the steps in Algorithm 3 as follows:

- Step 0: Initialization, given ϵ, μ and \mathbf{x}_0 , let $t = 0, \mathbf{x}(0) = \mathbf{x}_0, y_k(0) = \sigma_k(\mathbf{x}_{S_k}(0))$; arbitrarily choose $\lambda(0)$.

- Step 1: Find $\mathbf{x}^* \in K^{(1)}$ that solves

$$(\mathbf{x}' - \mathbf{x}^*)^T \left[f(\mathbf{x}^*) - A^T [\boldsymbol{\lambda}(t) - H(A\mathbf{x}^* + \mathbf{y}(t))] \right] \geq 0, \quad (17)$$

for $\forall \mathbf{x}' \in K$, and set $\mathbf{x}(t+1) = \mathbf{x}^*$.

- Step 2: Find $\mathbf{y}^* \in K^{(2)}$ that solves

$$(\mathbf{y}' - \mathbf{y}^*)^T \left[f(\mathbf{x}^*) - [\boldsymbol{\lambda}(t) - H(A\mathbf{x}(t+1) + \mathbf{y}^*)] \right] \geq 0, \quad (18)$$

for $\forall \mathbf{y}' \in K$, and set $\mathbf{y}(t+1) = \mathbf{y}^*$.

- Step 3: Set

$$\boldsymbol{\lambda}(t+1) = \boldsymbol{\lambda}(t) - H(A\mathbf{x}(t+1) - \mathbf{y}(t+1)) \quad (19)$$

- Step 4: Convergence verification: If $\|\boldsymbol{\omega}(t+1) - \boldsymbol{\omega}(t)\|_\infty < \epsilon$, then stop. Otherwise let $t \leftarrow t+1$ and go back to Step 1.

When we have $\mathbf{y}(t+1) = \mathbf{y}(t)$ and $\boldsymbol{\lambda}(t+1) = \boldsymbol{\lambda}(t)$, $\boldsymbol{\omega}(t+1) = (\mathbf{x}(t+1), \mathbf{y}(t+1), \boldsymbol{\lambda}(t+1))$ is the solution to our SVI₂. We denote the unique solution as $\boldsymbol{\omega}^* = (\mathbf{x}^*, \mathbf{y}^*, \boldsymbol{\lambda}^*)$. From Eqn (18) and (19), we have the following from Section 2 of (He, 2009),

$$\begin{aligned} & \|\mathbf{y}(t+1) - \mathbf{y}^*\|_H^2 + \|\boldsymbol{\lambda}(t+1) - \boldsymbol{\lambda}^*\|_{H^{-1}}^2 \\ & \leq \left(\|\mathbf{y}(t) - \mathbf{y}^*\|_H^2 + \|\boldsymbol{\lambda}(t) - \boldsymbol{\lambda}^*\|_{H^{-1}}^2 \right) \\ & \quad - \left(\|\mathbf{y}(t+1) - \mathbf{y}(t)\|_H^2 + \|\boldsymbol{\lambda}(t+1) - \boldsymbol{\lambda}(t)\|_{H^{-1}}^2 \right) \\ & < \|\mathbf{y}(t) - \mathbf{y}^*\|_H^2 + \|\boldsymbol{\lambda}(t) - \boldsymbol{\lambda}^*\|_{H^{-1}}^2, \end{aligned} \quad (20)$$

which shows the contraction property of the sequence $\{(\mathbf{y}(t), \boldsymbol{\lambda}(t))\}$ and thus proves the convergence of the algorithm.

A more detailed proof of convergence of the above steps in Eqn (17)-(19) is covered in (Gabay and Mercier, 1976; Glowinski and Oden, 1985), and a more generalized version of the above steps and convergence proofs are covered in (Tseng, 1990; Lions and Mercier, 1979). \square

D Proof of Theorem 3

D.1 Full version of HH-BRD

We will first show the full version of HH-BRD, suppose the superlevel partitions is taken between level $q-1$ and level q , then for $i = 1, \dots, N^{(1)}$,

$$\begin{aligned} \hat{u}_i^{(sl_1)} &= \sum_{l=1}^{q-1} u_{k_{il}}^{(l)}(x_{k_{il}}, \mathbf{x}_{I_{k_{il}}}) \\ & \quad - L_{k_{iq}}^{(sl_1, sl_2)} \left(\sigma_{k_{iq}}^{(1,q)}(\mathbf{x}_{S_{k_{iq}}^{(1,q)}}, x_{k_{iq}}^{(q)}) \right), \end{aligned} \quad (21)$$

where

$$S_p^{(1,q)} = \{a_i^{(1)} \mid k_{iq} = p\},$$

$$\sigma_p^{(1,q)}(\mathbf{x}_{S_p^{(1,q)}}) = \sum_{a_i^{(1)} \in S_p^{(1,q)}} x_i^{(1)}.$$

And for $j = 1, \dots, N^{(q)}$

$$\begin{aligned} \hat{u}_j^{(sl_2)} &= \sum_{l=q}^L u_{k_{jl}}^{(l)}(x_{k_{jl}}, \mathbf{x}_{I_{k_{jl}}}) \\ & \quad - L_j^{(sl_2, sl_1)} \left(x_j^{(q)}, \sigma_j^{(1,q)}(\mathbf{x}_{S_j^{(1,q)}}) \right). \end{aligned} \quad (22)$$

Please refer to Algorithm 5 for the pseudo code of the full version of this algorithm. The loss function updates are similar to that of Algorithm 3.

ALGORITHM 5: Hybrid Hierarchical BRD(Full Version)

Initialize the game, $t = 0, x_i^{(1)}(0) = (\mathbf{x}_0)_i, i = 1, \dots, N^{(0)}$

for $l = 2:L$ **do**

for $k = 1:N^{(l)}$ **do**

$\mathbf{x}_k^{(l)}(0) = \sigma_k^{(l)}(\mathbf{x}_{S_k^{(l)}}(0));$

end

end

while *not converged* **do**

for $k = 1:N^{(q)}$ (*Meta-Level-1 Penalty Update*) **do**

 Update $L_k^{(sl_1, sl_2)}$

end

for $i = 1 : N^{(1)}$ (*Level-1/Meta-Level-1 Gaming*) **do**

$x_i^{(1)}(t+1) =$

$$BR_i \left(\mathbf{x}_{I_i}^{(1)}(t), \mathbf{x}_{I_{k_{i2}}}^{(2)}(t), \dots, \mathbf{x}_{I_{k_{iq}}}^{(q)}(t), \hat{u}_i^{(sl_1)} \right)$$

end

for $l = 2:q-1$ (*Level-2 to Level- q Aggregation*) **do**

for $j = 1:N^{(l)}$ **do**

$\mathbf{x}_j^{(l)}(t+1) = \sigma_j^{(l)}(\mathbf{x}_{S_j^{(l)}}(t+1))$

end

end

for $k = 1:N^{(q)}$ (*Meta-Level-2 Penalty Update*) **do**

 Update $L_k^{(sl_2, sl_1)}$

end

for $j = 1 : N^{(q)}$ (*Level- q /Meta-Level-2 Gaming*) **do**

$$x_j^{(q)}(t+1)$$

$$= BR_j \left(\sigma_j^{(1,q)}(\mathbf{x}_{S_j^{(1,q)}}), \mathbf{x}_{I_j}^{(q)}(t), \mathbf{x}_{I_{k_{j(q+1)}}}^{(q+1)}(t), \dots, \right.$$

$$\left. \mathbf{x}_{I_{k_{jL}}}^{(L)}(t), \hat{u}_j^{(sl_2)} \right)$$

end

for $l = q+1:L$ (*Level-2 to Level- $q+1$ Aggregation*) **do**

for $p = 1:N^{(l)}$ **do**

$\mathbf{x}_p^{(l)}(t+1) = \sigma_p^{(l)}(\mathbf{x}_{S_p^{(l)}}(t+1))$

end

end

$t \leftarrow t+1;$

end

D.2 Proof of Theorem

We will first construct an equivalent 2-level game to the L -level game where $L > 2$, and then show that the action profile update trajectories are the same for the original game and the equivalent game. Finally, the convergence of the equivalent game follows Theorem 2 and thus Algorithm 4 guarantees convergence.

Proof. We define the following counter-part for utility component $u_i^{(l)}(x_i^{(l)}, \mathbf{x}_{I_i}^{(l)})$ ($1 < l < q$)

$$\underline{u}_i^{(l)}(\mathbf{x}_{S_i^{(1,l)}}, \mathbf{x}_{S_i^{(1,l)}}) = u_i^{(l)}(x_i^{(l)}, \mathbf{x}_{I_i}^{(l)}), \quad (23)$$

when $x_i^{(l)} = \sigma_i^{(1,l)}(\mathbf{x}_{S_i^{(1,l)}}), \forall i, \forall l \in \{2, \dots, q-1\}$. Both $\mathbf{x}_{S_i^{(1,l)}}$ and $\mathbf{x}_{S_i^{(1,l)}}$ are level-1 action profiles. This is exactly how we create the utility functions under the flat perspective, where we expand the higher level aggregate actions down to level-1.

Similarly, we define the following counter-part for utility component $u_j^{(l)}(x_j^{(l)}, \mathbf{x}_{I_j}^{(l)})$ ($q < l \leq L$)

$$\underline{u}_j^{(l)}(\mathbf{x}_{S_j^{(q,l)}}, \mathbf{x}_{S_j^{(q,l)}}) = u_j^{(l)}(x_j^{(l)}, \mathbf{x}_{I_j}^{(l)}), \quad (24)$$

when $x_j^{(l)} = \sigma_j^{(q,l)}(\mathbf{x}_{S_j^{(q,l)}}), \forall j, \forall l \in \{q, \dots, L\}$. Both $\mathbf{x}_{S_i^{(1,l)}}$ and $\mathbf{x}_{S_i^{(1,l)}}$ are level- q action profiles. This time we expand the higher level aggregate actions down to level- q instead of level-1.

So then we can define a “flattened” super-level-1 utility function counterpart for $u_i^{(sl_1)}$ as follows

$$\begin{aligned} \bar{u}_i^{(sl_1)}(x_i^{(1)}, x_{I_i}^{(1)}) &= \sum_{l=1}^{q-1} \underline{u}_{k_{il}}^{(l)}(\mathbf{x}_{S_{k_{il}}^{(1,l)}}, \mathbf{x}_{S_{I_{k_{il}}}^{(1,l)}}) \\ &\quad - L_{k_{iq}}^{(sl_1, sl_2)} \left(\sigma_{k_{iq}}^{(1,q)}(\mathbf{x}_{S_{k_{iq}}^{(1,q)}}), x_{k_{iq}}^{(q)} \right), \end{aligned} \quad (25)$$

where

$$\underline{I}_i^{(sl_1)} = \{a_j^{(1)} | k_{jq} = k_{iq}, j \neq i\}.$$

Similarly, for meta-level 2, we can define a “flattened”(to level- q) function counterpart for $u_j^{(sl_2)}$ as follows

$$\begin{aligned} \bar{u}_j^{(sl_2)}(x_j^{(q)}, x_{I_j}^{(sl_2)}) &= \sum_{l=q}^L \underline{u}_{k_{jl}}^{(l)}(\mathbf{x}_{S_{k_{jl}}^{(q,l)}}, \mathbf{x}_{S_{I_{k_{jl}}}^{(q,l)}}) \\ &\quad - L_j^{(sl_2, sl_1)} \left(x_j^{(q)}, \sigma_j^{(1,q)}(\mathbf{x}_{S_j^{(1,q)}}) \right), \end{aligned} \quad (26)$$

where

$$\underline{I}_j^{(sl_2)} = \{a_p^{(q)} | k_{pL} = k_{jL}, p \neq j\}.$$

So now we can create a 2-level game where the level-1 (resp. level- q) agents in the original game become the level-1 (resp. level-2) agents in the new game with utility functions defined in Eqn (25) (resp. Eqn (26)). Based on Theorem 2,

we know that if we apply SH-BRD, we can converge to the unique NE of the game under Assumptions 1-3.

Then it remains to show that given the same initialization, applying HH-BRD in the original game and the MS-BRD in the new 2-level game generate the same level-1 action profile update trajectory. This can be shown using induction.

We know from initialization that

$$x_i^{(l)}(0) = \sigma_i^{(1,l)}(\mathbf{x}_{S_i^{(1,l)}}(0)), \forall i, \forall l \in \{2, \dots, q-1\},$$

$$x_j^{(l)}(0) = \sigma_j^{(q,l)}(\mathbf{x}_{S_j^{(q,l)}}(0)), \forall j, \forall l \in \{q, \dots, L\}.$$

Then based on Eqn (23), we know that

$$\begin{aligned} &u_i^{(sl_1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}(0), \dots, x_{k_{iq}}^{(q)}(0)) \\ &= \bar{u}_i^{(sl_1)}(x_i^{(1)}, x_{I_i}^{(1)}(0)) \\ \Leftrightarrow &BR_i(\mathbf{x}_{I_i}^{(1)}(0), \dots, x_{k_{iq}}^{(q)}(0), u_i^{(sl_1)}) \\ &= BR_i(x_{I_i}^{(1)}(0), \bar{u}_i^{(sl_1)}), \end{aligned}$$

and thus when $t = 1$, $\mathbf{x}^{(1)}(t)$ are the same when applying HH-BRD in the original game and the MS-BRD in the new 2-level game. Similarly, $\mathbf{x}^{(q)}(1)$ are the same based on Eqn (24).

Suppose $\mathbf{x}^{(1)}(t)$ and $\mathbf{x}^{(q)}(t)$ are the same for the two dynamics for $t = 0, 1, \dots, T$, we need to show that $\mathbf{x}^{(1)}(t)$ and $\mathbf{x}^{(q)}(t)$ are the same for $t = T + 1$ to complete the proof.

Again, based on Eqn (23), we know that

$$\begin{aligned} &u_i^{(sl_1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}(T), \dots, x_{k_{iq}}^{(q)}(T)) \\ &= \bar{u}_i^{(sl_1)}(x_i^{(1)}, x_{I_i}^{(1)}(T)) \\ \Leftrightarrow &BR_i(\mathbf{x}_{I_i}^{(1)}(T), \dots, x_{k_{iq}}^{(q)}(T), u_i^{(sl_1)}) \\ &= BR_i(x_{I_i}^{(1)}(T), \bar{u}_i^{(sl_1)}), \end{aligned}$$

which implies $\mathbf{x}^{(1)}(T+1)$ are the same for the two dynamics and similarly $\mathbf{x}^{(q)}(T+1)$ are the same based on Eqn (24). \square

E Data Generation for Numerical Experiments

We introduce the data generation procedures for both games with linear best response and non-linear best response in this part.

First of all, for both type of games, we create an adjacency matrix for each of the groups on every level. This matrix has 0 diagonal elements and for the off-diagonal elements, the existence of a directed edge subjects to the Bernoulli distribution where there is a fixed P_{exist} . Then if a directed edge exist, the edge weight is generated by choosing a value from $[0, 1]$ uniformly at random. Later, we will multiply these matrices with different scalars to adjust the values so that Assumption 3 holds. These matrices have 0 diagonal elements because they capture the dependencies of agents on each other, or equivalently, they are used to model the external impact the agents receive from the network. The internal impact are modeled by cost functions and marginal benefit terms that only depend on an agent’s own action.

E.1 Linear Best Response Games

For games with linear best response, we generated a 2-level game with 100 groups and 10,000 level-1 agents. The adjacency matrix generation follows $P_{exist} = 0.1$, which creates a rather sparse network. Each level-2 group $S_k^{(2)}$ contains 100 members, and we use W_k to denote the corresponding adjacency matrix. We use V to denote the level-2 adjacency matrix. From Eqn (6), we know that for each level-1 agent, the utility function is

$$u_i(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}, \mathbf{x}_{I_{k_{i2}}}^{(2)}) = u_i^{(1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}) + u_{k_{i2}}^{(2)}(x_{k_{i2}}^{(2)}, \mathbf{x}_{I_{k_{i2}}}^{(2)}),$$

where

$$u_i^{(1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}) = b_i x_i^{(1)} + x_i^{(1)} \left(\sum_{j \in I_i} (W_{k_{i2}})_{r_i r_j} x_j^{(1)} \right) - c_i (x_i^{(1)})^2,$$

$$u_{k_{i2}}^{(2)}(x_{k_{i2}}^{(2)}, \mathbf{x}_{I_{k_{i2}}}^{(2)}) = x_{k_{i2}}^{(2)} \left(\sum_{p \neq k} V_{kp} x_p^{(2)} \right).$$

We choose the cost coefficients c_i to be large enough so that the $\Upsilon(F)$ satisfies the P_Υ condition (from Appendix A, strong diagonal dominance implies P_Υ condition). In the experiments, the $\rho(\Gamma)$ (See Appendix A for Γ) has a value between $[0.7, 0.8]$.

Then under the flat perspective, a level-1 agent $a_i^{(1)}$ has the following utility function

$$u_i^{flat}(x_i^{(1)}, x_{-i}^{(1)}) = b_i x_i^{(1)} + x_i^{(1)} \left(\sum_{j \neq i} W_{ij}^{flat} x_j^{(1)} \right) - c_i (x_i^{(1)})^2 + d_i,$$

where

$$d_i = \sum_{j \in I_i} x_j^{(1)} \left(\sum_{p \notin S_{k_{i2}}^{(2)}} W_{jp}^{flat} x_p^{(1)} \right),$$

$$W^{flat} = \begin{bmatrix} W_1 & V_{1,2} \cdot \mathbf{1} & \cdots & V_{1,100} \cdot \mathbf{1} \\ V_{2,1} \cdot \mathbf{1} & W_2 & \cdots & V_{2,100} \cdot \mathbf{1} \\ \vdots & \vdots & \ddots & \vdots \\ V_{100,1} \cdot \mathbf{1} & \cdots & V_{100,2} \cdot \mathbf{1} & W_{100} \end{bmatrix},$$

here $\mathbf{1}$ represents the all 1 matrix of suitable size (100×100).

E.2 General Best Response Games

For games with general (non-linear) best response, we generated data using the graphical game model similarly like the above. However, this time we use a mixed cost term that is a weighted sum of a quadratic component and an exponential component. Therefore, we can no longer represent the best response functions as linear functions and the best response computing now relies on gradient based optimization steps. In the experiments shown in the main article, the adjacency matrix is generated following $P_{exist} = 0.1$, which creates a sparse network. We also tried $P_{exist} = 1$ and the results on the dense networks are included in this part of the appendix.

We use $W_i^{(l)}$ to denote the adjacency matrix within $S_i^{(l)}$ and $W^{(L+1)}$ to denote the adjacency matrix between highest level agents. For the 2-level games with general best response, the utility components are set as follows

$$u_i^{(1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}) = b_i x_i^{(1)} + x_i^{(1)} \left(\sum_{j \in I_i} (W_{k_{i2}}^{(2)})_{r_i r_j} x_j^{(1)} \right) - c_i (x_i^{(1)})^2 - e^{0.1 x_i^{(1)}},$$

$$u_i^{(2)}(x_i^{(2)}, \mathbf{x}_{I_i}^{(2)}) = x_i^{(2)} \left(\sum_{j \neq i} (W_{k_{i3}}^{(3)})_{ij} x_j^{(2)} \right) - |S_i^{(2)}| \cdot e^{0.1 x_i^{(2)} / |S_i^{(2)}|}.$$

For 3-level games with general best response, the components in level-1 and 2 remain the same, and the level-3 components are

$$u_i^{(2)}(x_i^{(3)}, \mathbf{x}_{I_i}^{(3)}) = x_i^{(3)} \left(\sum_{j \neq i} W_{ij}^{(4)} x_j^{(3)} \right) - |S_i^{(1,3)}| \cdot e^{0.1 x_i^{(3)} / |S_i^{(1,3)}|}.$$

For the 2-level games with linear/nonlinear best response, the utility components are set as follows

$$u_i^{(1)}(x_i^{(1)}, \mathbf{x}_{I_i}^{(1)}) = b_i x_i^{(1)} + x_i^{(1)} \left(\sum_{j \in I_i} (W_{k_{i2}}^{(2)})_{r_i r_j} x_j^{(1)} \right) - c_i (x_i^{(1)})^2,$$

$$u_i^{(2)}(x_i^{(2)}, \mathbf{x}_{I_i}^{(2)}) = x_i^{(2)} \left(\sum_{j \neq i} (W_{k_{i3}}^{(3)})_{ij} x_j^{(2)} \right) - |S_i^{(2)}| \cdot e^{0.1 x_i^{(2)} / |S_i^{(2)}|}.$$

Again, the adjacency matrix and the cost terms will be scaled to ensure that Assumption 3 holds, and in the experiments, the $\rho(\Gamma)$ (See Appendix A for Γ) has a value between $[0.7, 0.8]$.

Hyperparameter settings: besides the parameters in the graphical games, the parameter $h_i^{(l)}$ in the loss function updates in Eqn (10) is chosen arbitrarily. These parameters can also be referred to as ‘‘penalty parameters’’. In our experiments, the performance over these parameters are rather smooth under assumption 3. The hyperparameters $h_i^{(l)}$ are set to the same value on each level l . In the 2-level case, we perform a binary search on these hyperparameter, where each value is tested for 5 runs to see the average performance. For the 3-level case, we need to determine 2 hyperparameter values, and this is done by a fixed step size search performed iteratively on the two values. We tune the first one, each value is tested for 5 runs like the above, while fixing the second value, after that, we switch to the tuning of the second value and this process keeps iteratively. The parameters we used in the numerical experiments are

- 2-Level game: $h_i^{(2)} = 0.2, 0.1, 0.06, 0.03, 0.01$; for network sizes $30^2, 50^2, 100^2, 200^2, 500^2$ respectively. With tuning range $[0, 0.5]$.

- 3-Level game:

For SH-BRD: $(h_i^{(2)}, h_j^{(3)}) = (0.65, 0.1), (0.32, 0.03), (0.2, 0.01), (0.12, 0.006), (0.04, 0.003)$; for network sizes $10^3, 20^3, 30^3, 50^3, 100^3$ respectively. With tuning range $[0, 0.5]^2$ and tuning step 0.002.

For HH-BRD: $h_i^{(sl_1)} = 0.7, 0.3, 0.21, 0.125, 0.063$ for network sizes $10^3, 20^3, 30^3, 50^3, 100^3$ respectively. With tuning range $[0, 0.5]$

Under the current parameter settings, we still haven't bring out the best performances of SH-BRD, and HH-BRD. In act, the performance gap between the current setting and the optimal setting won't be too large since the best response steps are well-posed. And even with their sub-optimal performances, we have seen their advantages over other algorithms.

In (He, Yang, and Wang, 2000), the authors mentioned an adaptive method to generate the penalty parameter matrix H which is generally not diagonal, that can speed up the problem solving steps. This will be an interesting direction to generalize our current algorithm when the best response functions become more ill-posed in the future.

E.3 CPU Specs:

- CPU: 6 cores, 12 threads, 2.60/4.50 GHz, 12MB Cache
- OS: Windows 10
- Software: Python 3.7
- RAM: 16 GB

E.4 Results on Dense Networks

Size	BRD	MS-BRD	SH-BRD
30^2	$(2.97 \pm 0.24) \times 10^7$	$(9.91 \pm 0.81) \times 10^5$	$(8.31 \pm 0.66) \times 10^5$
50^2	$(2.41 \pm 0.22) \times 10^8$	$(4.83 \pm 0.45) \times 10^6$	$(3.27 \pm 0.30) \times 10^6$
100^2	$(4.07 \pm 0.34) \times 10^9$	$(4.07 \pm 0.34) \times 10^7$	$(3.04 \pm 0.22) \times 10^7$
200^2	$(6.66 \pm 0.62) \times 10^{10}$	$(3.33 \pm 0.31) \times 10^8$	$(2.44 \pm 0.17) \times 10^8$
500^2	$(2.72 \pm 0.29) \times 10^{12}$	$(5.53 \pm 0.49) \times 10^9$	$(3.26 \pm 0.26) \times 10^9$

Table 5: Convergence and complexity (flops) comparison with linear best response under multiple initialization, dense network.

We can see that though the results in linear best response games are very different in sparse and dense networks, the results in games with non-linear best responses are quite similar in both types of networks. In games with linear best responses, the standard deviation results from different initialization. For the same game, one initial action profile's distance (measured in Euclidean norm) to the equilibrium point can be 20 times to the distance of another initial action profile. This results in different number of iterations of the algorithm before convergence. However, it only takes about 20% more iterations for a "distant" initial action profile to reach convergence, which shows that these algorithms

Size	BRD	MS-BRD	SH-BRD
30^2	0.99 ± 0.03	0.49 ± 0.02	0.24 ± 0.01
50^2	22.80 ± 0.05	1.83 ± 0.06	0.69 ± 0.01
100^2	1351 ± 7	13.28 ± 0.26	4.70 ± 0.06
200^2	> 18000	159.9 ± 0.8	58.07 ± 0.42
500^2	nan	3505 ± 54	1286 ± 20

Table 6: CPU times on a single machine on 2-Level games with general best response functions, dense network; All times are in seconds.

Size	BRD	MS-BRD	SH-BRD
30^2	1.63 ± 0.12	0.57 ± 0.02	0.028 ± 0.002
50^2	30.65 ± 0.35	1.94 ± 0.03	0.051 ± 0.003
100^2	1660 ± 3	13.93 ± 0.25	0.33 ± 0.02
200^2	> 18000	163.1 ± 1.4	1.32 ± 0.04
500^2	nan	3416 ± 52	29.37 ± 0.91

Table 7: CPU times on a single machine for 2-Level, linear/nonlinear best-response games, dense network; All times are in seconds.

Size	BRD	MS-BRD	SH-BRD	HH-BRD
10^3	1.25 ± 0.02	0.39 ± 0.01	0.57 ± 0.02	0.34 ± 0.01
20^3	617.3 ± 4.7	2.85 ± 0.07	4.50 ± 0.06	2.56 ± 0.06
30^3	> 18000	10.25 ± 0.25	17.87 ± 0.14	9.53 ± 0.09
50^3	nan	58.04 ± 0.32	100.8 ± 0.41	51.86 ± 0.24
100^3	nan	926.8 ± 6.4	2131 ± 11	780.9 ± 3.0

Table 8: CPU times in seconds on a single machine on 3-Level, general best response games, dense network; All times are in seconds.

have good convergence property under Assumptions 1-3. In games with non-linear best responses, the standard deviations of CPU times are relatively small(around 1%) compared to the mean values, and it shows that the performance of all algorithms are stable with a fixed initial action profile.

F Algorithm Performances and Network Sizes

In this part, we present some results that show the algorithms' performances with different network sizes in 2-level games.

Figure 2 shows the number of flops per iteration for the three algorithms in $I \times M$ games where I is the number of agents in each group and M the number of groups in the network. Both Algorithms 2 and 4 outperform Algorithm 1. Algorithm 4 generally has lower complexity per iteration compared to Algorithm 2 since it has less input in every sub-problem and the number of sub-problems are similar in Algorithm 2 and 4 when the group sizes are large. However, when group sizes are small compared to the number of groups, Algorithm 2 and 4 are similar per iteration.




Figure 2: Complexity per iteration for linear best response.

G Reverse Engineer Multi-scale Structure

A question that naturally arises is whether sparsity in the network can be exploited when the multi-scale structure is not readily available. The utility function in Eqn (6) suggests that such reverse engineering is possible if the game satisfies:

1. An agent is either connected to all agents in another group or not connected to any agent in that group; If so, we can create a set of possible group partitions.
2. Based on the partition in the previous step, agents in one group have the same dependency on an agent in another group.
3. Based on the partition, we can represent the groups' aggregate actions from their members' actions using some aggregate functions.
4. Based on the partition, the original utility function of each agent can be separated to components on different levels, each component only based on the actions and dependencies on the corresponding level.

An example of the first condition is shown in Figs. 3 and 4. For the other conditions, the "flattened" utility functions used in Appendix E are good examples.




Figure 3: Ungrouped.




Figure 4: Grouped.

H Flow Charts of the Algorithms

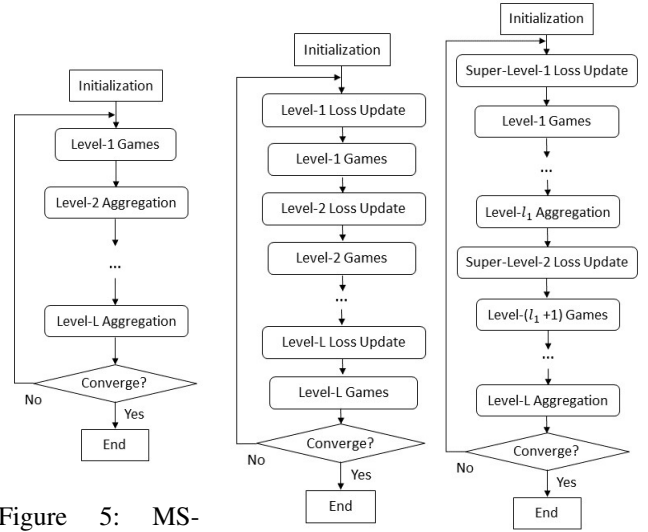


Figure 5: MS-BRD

Figure 6: SH-BRD Figure 7: HH-BRD