

# SOLVING STRUCTURED HIERARCHICAL GAMES USING DIFFERENTIAL BACKWARD INDUCTION

Zun Li<sup>1</sup>, Feiran Jia<sup>2</sup>, Aditya Mate<sup>3</sup>, Shahin Jabbari<sup>4</sup>, Mithun Chakraborty<sup>1</sup>,  
Milind Tambe<sup>3</sup>, and Yevgeniy Vorobeychik<sup>5</sup>

<sup>1</sup>University of Michigan, Ann Arbor

<sup>2</sup>Pennsylvania State University

<sup>3</sup>Harvard University

<sup>4</sup>Drexel University

<sup>5</sup>Washington University in St. Louis

{lizun, dcsmc}@umich.edu

fzj5059@psu.edu

{aditya.mate, milind.tambe}@g.harvard.edu

shahin@drexel.edu

yvorobeychik@wustl.edu

## ABSTRACT

From large-scale organizations to decentralized political systems, hierarchical strategic decision making is commonplace. We introduce a novel class of *structured hierarchical games (SHGs)* that formally capture such hierarchical strategic interactions. In an SHG, each player is a node in a tree, and strategic choices of players are sequenced from root to leaves, with root moving first, followed by its children, then followed by their children, and so on until the leaves. A player’s utility in an SHG depends on its own decision, and on the choices of its parent and *all* the tree leaves. SHGs thus generalize simultaneous-move games, as well as Stackelberg games with many followers. We leverage the structure of both the sequence of player moves as well as payoff dependence to develop a novel gradient-based back propagation-style algorithm, which we call *Differential Backward Induction (DBI)*, for approximating equilibria of SHGs. We provide a sufficient condition for convergence of DBI and demonstrate its efficacy in finding approximate equilibrium solutions to several SHG models of hierarchical policy-making problems.

## 1 INTRODUCTION

The COVID-19 pandemic has revealed considerable strategic tension among the many parties involved in decentralized hierarchical policy-making. For example, recommendations by the World Health Organization are sometimes heeded, and other times discarded by nations, while subnational units, such as provinces and urban areas, may in turn take a policy stance (such as on lockdowns, mask mandates, or vaccination priorities) that is not congruent with national policies. Similarly, in the US, policy recommendations at the federal level can be implemented in a variety of ways by the states, while counties and cities, in turn, may comply with state-level policies, or not, potentially triggering litigation Hill & Varone (2021). Central to all these cases is that, besides this strategic drama, what ultimately determines infection spread is how policies are implemented *at the lowest level*, such as by cities and towns, or even individuals. Similar strategic encounters routinely play out in large-scale organizations, where actions throughout the management hierarchy are ultimately reflected in the decisions made at the lowest level (e.g., by the employees who are ultimately involved in production), and these lowest-level decisions play a decisive role in the organizational welfare.

We propose a novel model of hierarchical decision making which is a natural stylized representation of strategic interactions of this kind. Our model, which we term *structured hierarchical games (SHGs)*, represents each player by a node in a tree hierarchy. The tree plays two roles in SHGs. First, it captures the sequence of moves by the players: the root (the lone member of level 1 of the hierarchy) makes the first strategic choice, its children (i.e., all nodes in level 2) observe the root’s choice and

follow, their children then follow in turn, and so on, until we reach the leaf node players who move upon observing their predecessors’ choices. Second, the tree partially captures strategic dependence: a player’s utility depends on its own strategy, that of its parent, and the strategies of *all of the leaf nodes*. The sequence of moves in our model naturally captures the typical sequence of decisions in hierarchical policy-making settings, as well as in large organizations, while the utility structure captures the decisive role of leaf nodes (e.g., individual compliance with vaccination policies), as well as hierarchical dependence (e.g., employee dependence on a manager’s approval of their performance, or state dependence on federal funding). Significantly, the *SHG* model generalizes a number of well-established models of strategic encounters, including (a) simultaneous-move games (captured by a 2-level SHG with the root having a single dummy action), (b) Stackelberg (leader-follower) games (a 2-level game with a single leaf node) Von Stackelberg (1952); Fiez et al. (2020), and (c) single-leader multi-follower Stackelberg games (e.g., a Stackelberg security game with a single defender and many attackers) Basilico et al. (2016); Coniglio et al. (2020).

Our second contribution is a novel gradient-based algorithm for approximately computing subgame-perfect equilibria of *SHGs*. Specifically, we propose *Differential Backward Induction (DBI)*, which is a backpropagation-style gradient ascent algorithm that leverages both the sequential structure of the game, as well as the utility structure of the players. As *DBI* involves simultaneous gradient updates of players in the same level (particularly at the leaves), convergence is not guaranteed in general (as is also the case for best-response dynamics Fudenberg et al. (1998)). Viewing *DBI* as a dynamical system, we provide a sufficient condition for its convergence to a stable point. Our results also imply that in the special case of two-player zero-sum Stackelberg games, *DBI* converges to a local Stackelberg equilibrium Fiez et al. (2020); Wang et al. (2020).

Finally, we demonstrate the efficacy of *DBI* in finding approximate equilibrium solutions to several classes of *SHGs*. First, we use *DBI* to solve a recently proposed game-theoretic model of 3-level hierarchical epidemic policy making. Second, we apply *DBI* to solve a hierarchical variant of a public goods game, which naturally captures the decentralization of decision making in public good investment decisions, such as investments in sustainable energy. Third, we evaluate *DBI* in the context of a hierarchical security investment game, where hierarchical decentralization (e.g., involving federal government, industry sectors, and particular organizations) can also play a crucial role. In all of these, we show that *DBI* significantly outperforms the state of the art approaches that can be applied to solve games with hierarchical structure.

**Related Work** *SHGs* generalize both simultaneous-move games and Stackelberg games with multiple followers (Leyffer & Munson, 2005; Basilico et al., 2016). They are also related to *graphical games* Kearns et al. (2001) in capturing utility dependence structure, although *SHGs* also capture sequential structure of decisions. Several prior approaches use gradient-based methods for solving games with particular structure. A prominent example is generative adversarial networks (GANs), though these are zero-sum games (Goodfellow et al., 2014; Jin et al., 2020; Nagarajan & Kolter, 2017; Daskalakis & Panageas, 2018; Mertikopoulos & Zhou, 2019; Mescheder et al., 2017). Ideas from learning GANs have been adopted in gradient-based approaches to solve multi-player general-sum games (Mazumdar et al., 2020; Balduzzi et al., 2018; Chasnov et al., 2020; Ibrahim et al., 2020; Letcher, 2021; Lin et al., 2020; Mertikopoulos & Zhou, 2019). However, all of these approaches assume a simultaneous-move game. A closely-related thread to our work considers gradient-based methods for bi-level optimization (Li & Başar, 1987; Shaban et al., 2019). Several related efforts consider gradient-based learning in Stackelberg games, and also use the implicit function theorem to derive gradient updates (Amin et al., 2016; Fiez et al., 2020; Nguyen et al., 2020; Wang et al., 2020; 2021). We significantly generalize these ideas by considering an arbitrary hierarchical game structure.

Jia et al. (2021) recently considered a stylized 3-level SHG for pandemic policy making, and proposed several non-gradient-based algorithms for this problem. We compare with their approach in Section 4.

## 2 STRUCTURED HIERARCHICAL GAMES

**Notation** We use bold lower-case letters to denote vectors. Let  $f$  be a function of the form  $f(\mathbf{x}, \mathbf{y}) : \mathbb{R}^d \times \mathbb{R}^{d'} \rightarrow \mathbb{R}^{d''}$ . We use  $\nabla_{\mathbf{x}} f$  to denote the partial derivative of  $f$  with respect to  $\mathbf{x}$ . When there is functional dependency between  $\mathbf{x}$  and  $\mathbf{y}$ , we use  $D_{\mathbf{x}} f$  to denote the total derivative of  $f(\mathbf{x}, \mathbf{y}(\mathbf{x}))$  with respect to  $\mathbf{x}$ . We use  $\nabla_{\mathbf{x}, \mathbf{x}}^2 f$  and  $\nabla_{\mathbf{x}, \mathbf{y}}^2 f$  to denote the second-order partial derivatives

and  $D_{\mathbf{x}, \mathbf{x}}^2 f$  to denote the second-order total derivative of  $f$ . For a mapping  $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , we use  $f^t(\mathbf{x})$  to denote  $t$  iterative applications of  $f$  on  $\mathbf{x}$ . For mappings  $f_1 : \mathbb{R}^d \rightarrow \mathbb{R}^d$  and  $f_2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ , we define  $(f_1 \circ f_2)(\mathbf{x}) \triangleq f_1(f_2(\mathbf{x}))$  and  $(f_1 + f_2)(\mathbf{x}) \triangleq f_1(\mathbf{x}) + f_2(\mathbf{x})$ . Moreover, for a given  $\epsilon \in \mathbb{R}^{\geq 0}$  and  $\mathbf{x} \in \mathbb{R}^d$ , we define the  $\epsilon$ -ball around  $\mathbf{x}$  as  $\mathbb{B}_\epsilon(\mathbf{x}) = \{\mathbf{x}' \in \mathbb{R}^d \mid \|\mathbf{x} - \mathbf{x}'\|_2 < \epsilon\}$ . Finally,  $\mathbf{I}$  denotes an identity matrix.

**Formal Model** A structured hierarchical game (SHG)  $\mathcal{G}$  consists of the set  $\mathcal{N}$  of  $n$  players. Each player  $i$  is associated with a set of actions  $\mathcal{X}_i \subseteq \mathbb{R}^{d_i}$ . The players are partitioned across  $L$  levels, where  $\mathcal{N}_l$  is the set of  $n_l$  players occupying level  $l$ . Let  $l_i$  denote the level occupied by player  $i$ . This *hierarchical* structure of the game is illustrated in Figure 1 where players correspond to nodes and levels are marked by dashed boundaries. The hierarchy plays two crucial roles: 1) it determines the order of moves, and 2) it partly determines utility dependence among players. Specifically, the temporal pattern of actions is as follows: level 1 has a single player, the *root*, who chooses an action first, followed by all players in level 2 making simultaneous choices, followed in turn by players in level 3, and so on until the *leaves* in the final level  $L$ . Players of level  $l$  only observe the actions chosen by all players of levels  $1, 2, \dots, l-1$ , but not their peers in the same level.

So, for example, pandemic social distancing and vaccination policies in the US are initiated by the federal government (including the Centers for Disease Control and Prevention), with states subsequently instituting their own policies, counties reacting to these by determining their own, and behavior of people ultimately influenced, but not determined, by the guidelines and enforcement policies by the local county/city.

Next, we describe the utility structure of the game as entailed by the SHG hierarchy. Each player  $i$  in level  $l_i > 1$  (i.e., any node other than the root) has a *unique parent* in level  $l_i - 1$ ; we denote the parent of node  $i$  by  $\text{PA}(i)$ . A player’s utility function is determined by 1) its own action, 2) the action of its parent, and 3) the actions of *all* players in level  $L$  (i.e., all *leaf* players). To formalize, let  $\mathbf{x}_l$  denote the joint action profile of all players in level  $l$ . Player  $i$ ’s utility function then has the form  $u_i(x_i, \mathbf{x}_L)$  if  $l_i = 1$ ,  $u_i(x_i, x_{\text{PA}(i)}, \mathbf{x}_L)$  if  $1 < l_i < L$ , and  $u_i(x_i, x_{\text{PA}(i)}, \mathbf{x}_{L,-i})$  if  $l_i = L$ , where  $\mathbf{x}_{L,-i}$  is the action profile of all players in level  $L$  other than  $i$ . For example, in our running pandemic policy example, the utility of a county depends on both the policy and enforcement strategy of its state (its *parent*) and on the ultimate pandemic spread and economic impact within it, both determined largely by the behavior of the county residents (*leaf nodes*). Note the considerable generality of the SHG model. For example, an arbitrary simultaneous-move game is a SHG with 2 levels and a “dummy” root node (utilities of all leaves depend on one another’s actions), and an arbitrary Stackelberg game (e.g., Stackelberg security game), even with many followers, can be modeled as a 2-level SHG with the leader as root and followers as leaves. Furthermore, while we have defined SHGs with respect to real-vector player action sets, it is straightforward to represent mixed strategies of finite-action games in this way by simply using a softmax function to map an arbitrary real vector into a valid mixed strategy.

**Solution Concept** Since an SHG has important sequential structure, it is natural to consider the *subgame perfect equilibrium* (SPE) as the solution concept (Osborne & Rubinstein, 2004). Here, we focus on pure-strategy equilibria. To begin, we note that in SHGs, the strategies of players in any level  $l > 1$  are, in general, functions of the complete history of play in levels  $1, \dots, l-1$ , which we denote by  $h_{<l} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{l-1})$ . Formally, a (pure) strategy of a player  $i$  is denoted by  $s_i(h_{<l})$ , which deterministically maps an arbitrary history  $h_{<l}$  into an action  $x_i \in \mathcal{X}_i$ . A *Nash equilibrium* of an SHG is then a strategy profile  $\mathbf{s} = (s_1, \dots, s_i, \dots, s_n)$  such that for all  $i \in \mathcal{N}$ ,  $u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s'_i, \mathbf{s}_{-i})$  for all possible alternative strategies for  $i$ ,  $s'_i$ . Here, we denote the realized payoff of  $i$  from profile  $\mathbf{s}$  by  $u_i(s_i, \mathbf{s}_{-i})$ . Next, we define a *level- $l$ -subgame* given  $h_{<l}$  as an SHG that includes only players at levels  $\geq l$ , with actions chosen in levels  $< l$  fixed to  $h_{<l}$ . A strategy profile  $\mathbf{s}$  is a *subgame perfect equilibrium* of SHG if it is a Nash equilibrium of every level- $l$ -subgame of SHG for every  $l$  and history  $h_{<l}$ . We prove in appendix A that our definition of SPE is equivalent to the standard SPE in an extensive-form representation of SHG.

While in principle we can compute an SPE of an SHG using backward induction, this cannot be done directly (i.e., by complete enumeration of actions of all players) as actions are real vectors. Moreover, even discretizing actions is of little help, as the hierarchical nature of the game leads to exponential explosion of the search space. We now present a gradient-based approach for approximating SPE

along the equilibrium path in an SHG that leverages the game structure to derive backpropagation-style gradient updates.

### 3 DIFFERENTIAL BACKWARD INDUCTION

In this section, we describe our gradient-based algorithm, Differential Backward Induction (DBI), for approximating an SPE, and then analyze its convergence. Just as gradient ascent does not, in general, identify a globally optimal solution to a non-convex optimization problem, DBI in general yields a solution which only satisfies first-order conditions (see Section B.1 for further details). Moreover, we leverage the structure of the utility functions to focus computation on an SPE in which strategies of players are only a function of their immediate parents.<sup>1</sup> In this spirit, we define *local* best response functions  $\phi_i : \mathbb{R}^{d_{PA(i)}} \rightarrow \mathbb{R}^{d_i}$  mapping a player  $i$ 's parent's action  $x_{PA(i)}$  to  $i$ 's action  $x_i$ ; note that the notation  $\phi_i$  is distinct from  $s_i$  above for  $i$ 's strategy to emphasize the fact that  $\phi_i$  is only locally optimal. Now, suppose that a player  $i$  is in the last level  $L$ . Local optimality of  $\phi_i$  implies that if  $x_i = \phi_i(x_{PA(i)})$ , then  $\nabla_{x_i} u_i(x_i, x_{PA(i)}, x_{L,-i}) = 0$  and  $\nabla_{x_i, x_i}^2 u_i(x_i, x_{PA(i)}, x_{L,-i}) < 0$ .<sup>2</sup>

Let  $\phi_l$  denote the local best response for all the players in level  $l$  given the actions of all players in level  $l-1$ . We can compose these local best response functions to define the function  $\Phi_l := \phi_L \circ \phi_{L-1} \circ \dots \circ \phi_{l+1} : \mathbb{R}^{d_{n_l}} \rightarrow \mathbb{R}^{d_{n_L}}$  i.e., the local best response of players in the last level  $L$  given the actions of the players in level  $l$ .<sup>3</sup> Then for any player ( $i$ ) in level  $l_i < L$ ,  $D_{x_i} u_i(x_i, x_{PA(i)}, \Phi_l(\langle x_i, x_{l,-i} \rangle)) = 0$  and  $D_{x_i, x_i}^2 u_i(x_i, x_{PA(i)}, \Phi_l(\langle x_i, x_{l,-i} \rangle)) < 0$ , where  $D_{x_i}$  is the total derivative with respect to  $x_i$  (as  $\Phi_l(\langle x_i, x_{l,-i} \rangle)$  is also a function of  $x_i$ ). Note that the functions  $\phi$  and  $\Phi$  are *implicit*, capturing the functional dependencies between actions of players in different levels at the local equilibrium.

Throughout, we make the following standard assumption on the utility functions (Dontchev & Rockafellar, 2009; Wang et al., 2020).

**Assumption 1.** For any  $x_i \in \mathcal{X}_i$ , the second-order partial derivatives of the form  $\nabla_{x_i, x_i}^2 u_i$  are non-singular.

#### 3.1 ALGORITHM

The DBI algorithm works in a bottom-up manner, akin to back-propagation: for each level  $l$ , we compute the total derivatives (gradients) of the utility functions and local best response maps ( $\phi$ ,  $\Phi$ ) based on analytical expressions that we derive below. We then propagate this information up to level  $l-1$ , as it is used to compute gradients for that level, and so on until level 1. Algorithm 1 gives the full DBI algorithm. In this algorithm,  $\text{CHD}(i)$  denotes the set of children of player  $i$  (i.e., nodes in level  $l_i + 1$  for whom  $i$  is the parent).

<sup>1</sup>Note that while we cannot guarantee that an SPE exists in SHGs in general, let alone those possessing the assumed structure, we find experimentally that our approach often yields good SPE approximations.

<sup>2</sup>For simplicity, we omit degenerate cases where  $\nabla_{x_i, x_i}^2 u_i = 0$  and assume all local maxima are strict.

<sup>3</sup>Note that in particular  $\Phi_L = \phi_L$ .




Figure 1: Schematic representation of an SHG. The utility of player  $i$  can have direct functional dependence *only* on the joint action of *all* shaded players.

---

**Algorithm 1** Differential Backward Induction (DBI)
 

---

**Input:** An SHG instance  $\mathcal{G}$

**Parameters:** Learning rate  $\alpha$ , maximum number of iterations  $T$  for gradient update

**Output:** A strategy profile

Randomly initialize  $\mathbf{x}^0 = \langle \mathbf{x}_1^0, \dots, \mathbf{x}_L^0 \rangle$

**for**  $t = 1, 2, \dots, T$  **do**

**for**  $l = L, L - 1, \dots, 1$  **do**

**for**  $i = 1, 2, \dots, n_l$  **do**

**if**  $l = L$  **then**

        Back-propagate  $D_{x_i} \Phi_i = \mathbf{I}$  to  $\text{PA}(i)$

        Set  $x_i^t \leftarrow x_i^{t-1} + \alpha \nabla_{x_i} u_i$

**else**

        Compute  $\nabla_{x_i} u_i, \nabla_{\mathbf{x}_L} u_i$  at  $\mathbf{x}^{t-1}$

        Compute  $D_{x_i} \phi_j, \forall j \in \text{CHD}(i)$  (Eqn. equation 5)

        Compute  $D_{x_i} \Phi_l$  (Eqn. equation 4)

        Back-propagate  $D_{x_i} \Phi_l$  to  $\text{PA}(i)$

        Compute  $D_{x_i} u_i = \nabla_{x_i} u_i + \nabla_{\mathbf{x}_L} u_i D_{x_i} \Phi_l$

        Set  $x_i^t \leftarrow x_i^{t-1} + \alpha D_{x_i} u_i$

**Return**  $\mathbf{x}^T$

---

DBI works in a backward message-passing manner, comparable to back-propagation: after each player has computed its total derivative, it passes (back-propagates)  $D_{x_i} \Phi_l$  to its direct parent; this information is, in turn, used by the parent to compute its own total derivative, which is passed to its own parent, and so on.

Algorithm 1 takes the total derivatives as given. We now derive closed-form expressions for these. We start from the last level  $L$ . Given the actions of players in level  $L - 1$ , the total derivative of a player  $i \in \mathcal{N}_L$  with respect to  $x_i$  is

$$D_{x_i} u_i(x_i, x_{\text{PA}(i)}, \mathbf{x}_{L,-i}) = \nabla_{x_i} u_i. \quad (1)$$

For a player  $i$  in level  $L - 1$ , the total derivative (at a local best response) is

$$\begin{aligned} D_{x_i} u_i(x_i, x_{\text{PA}(i)}, \phi_L(\langle x_i, \mathbf{x}_{L-1,-i} \rangle)) \\ = \nabla_{x_i} u_i + (\nabla_{\mathbf{x}_L} u_i)(D_{x_i} \phi_L), \end{aligned} \quad (2)$$

where  $\nabla_{\mathbf{x}_L} u_i$  is a  $1 \times d_{n_L}$  vector and  $D_{x_i} \phi_L$  is a  $d_{n_L} \times d$  matrix. The technical challenge here is to derive the term  $D_{x_i} \phi_L$  for  $i \in \mathcal{N}_{L-1}$ . Recall that  $\phi_L$  is the vectorized concatenation of the  $\phi_j$  functions for  $j \in \mathcal{N}_L$ . Since the local best response strategy of a player in level  $L$  only depends on its parent in level  $L - 1$ , the only terms in  $\phi_L$  that depend on  $x_i$  are the actions of  $\text{CHD}(i)$  in level  $L$ .

Consequently, it suffices to derive  $D_{x_i} \phi_j$  for  $j \in \text{CHD}(i)$ . Note that for these players  $j$ ,  $\nabla_{x_j} u_j = 0$  (by local optimality of  $\phi_L$ ). We will use this first-order condition to derive the expression for the total derivative using the *implicit function theorem*.

**Theorem 1** (Implicit Function Theorem (IFT) (Dontchev & Rockafellar, 2009, Theorem 1B.1)). *Let  $f(\mathbf{x}_1, \mathbf{x}_2) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$  be a continuously differentiable function in a neighborhood of  $(\mathbf{x}_1^*, \mathbf{x}_2^*)$  such that  $f(\mathbf{x}_1^*, \mathbf{x}_2^*) = 0$ . Also suppose  $\nabla_{\mathbf{x}_2} f$ , the Jacobian of  $f$  with respect to  $\mathbf{x}_2$ , is non-singular at  $(\mathbf{x}_1^*, \mathbf{x}_2^*)$ . Then around a neighborhood of  $\mathbf{x}_1^*$ , we have a local diffeomorphism  $\mathbf{x}_2^*(\mathbf{x}_1) : \mathbb{R}^d \rightarrow \mathbb{R}^d$  such that  $D_{\mathbf{x}_1} \mathbf{x}_2 = -(\nabla_{\mathbf{x}_2} f)^{-1} \nabla_{\mathbf{x}_1} f$ .*

To use Theorem 1, we set  $f = \nabla_{x_j} u_j$  (which satisfies the conditions of Theorem 1 by Assumption 1),  $\mathbf{x}_1 = x_i$  and  $\mathbf{x}_2 = x_j$  (recall that  $j \in \text{CHD}(i)$ ). By IFT, there exists  $\phi_j(x_i)$  such that  $D_{x_i} \phi_j = -(\nabla_{x_j, x_j}^2 u_j)^{-1} \nabla_{x_j, x_i}^2 u_j$ . Define  $\nabla_j^2 := \nabla_{x_j, x_i}^2 u_j$ . Then

$$\begin{aligned} (\nabla_{\mathbf{x}_L} u_i)(D_{x_i} \phi_L) &= - \sum_{j \in \text{CHD}(i)} (\nabla_{x_j} u_i) D_{x_i} \phi_j \\ &= - \sum_{j \in \text{CHD}(i)} (\nabla_{x_j} u_i) (\nabla_{x_j, x_j}^2 u_j)^{-1} \nabla_j^2. \end{aligned}$$

Plugging this into Equation equation 2, we obtain

$$\begin{aligned} D_{x_i} u_i(x_i, x_{\text{PA}(i)}, \phi_L(\mathbf{x}_{L-1})) \\ = \nabla_{x_i} u_i - \sum_{j \in \text{CHD}(i)} (\nabla_{x_j} u_i) (\nabla_{x_j, x_j}^2 u_j)^{-1} \nabla_j^2. \end{aligned} \quad (3)$$

For a level  $l < L - 1$ , the total derivative of player  $i \in \mathcal{N}_l$  in a local best response is  $D_{x_i} u_i = \nabla_{x_i} u_i + (\nabla_{\mathbf{x}_L} u_i)(D_{x_i} \Phi_l)$ , where

$$\begin{aligned} D_{x_i} \Phi_l &= (D_{\mathbf{x}_{l+1}} \Phi_{l+1})(D_{x_i} \mathbf{x}_{l+1}) \\ &= \sum_{j \in \text{CHD}(i)} (D_{x_j} \Phi_{l+1})(D_{x_i} \phi_j). \end{aligned} \quad (4)$$

Applying IFT, we get

$$D_{x_i} \phi_j = -(\nabla_{x_j, x_j}^2 u_j)^{-1} \nabla_{x_j, x_i}^2 u_j, \quad (5)$$

for  $j \in \text{CHD}(i)$ . We can apply the above procedure recursively for  $D_{\mathbf{x}_{l+1}} \Phi_{l+1}$  to derive the total derivative for players  $i \in \mathcal{N}_l$  for  $l < L - 1$ :

$$\begin{aligned} D_{x_i} u_i &= \nabla_{x_i} u_i + \left( \sum_{j \in \text{LEAF}(i)} (-1)^{L-l} \nabla_{x_j} u_i \right. \\ &\quad \left. \prod_{\eta \in \text{PATH}(j \rightarrow i)} (\nabla_{x_\eta, x_\eta}^2 u_\eta)^{-1} \nabla_{x_\eta, \mathbf{x}_{\text{PA}(\eta)}}^2 u_\eta \right), \end{aligned} \quad (6)$$

where  $\text{PATH}(j \rightarrow i)$  is an ordered list of nodes (players) lying on the unique path from  $j$  to  $i$ , excluding  $i$ . Note that Equation equation 6 is a generalization of Equation equation 3 where the  $\text{PATH}$  only consists of the leaf player.

While the above derivation assumes the  $\phi$  and  $\Phi$  functions are local best responses, in our algorithm in each iteration we evaluate these functional expressions for the total derivatives *at the current joint action profile*. This significantly reduces computational complexity and ensures that Algorithm 1 satisfies the first-order conditions upon convergence.

### 3.2 CONVERGENCE ANALYSIS

We analyze the convergence properties of our algorithm in appendix B

## 4 EXPERIMENTS

In this section, we empirically investigate the behavior of DBI in games where we can verify global stability. All our code is written in python. We ran our experiments on an Intel(R) Core(TM) i9-9820X CPU @ 3.30GHz.

We evaluate the performance in terms of quality of equilibrium approximation as a function of the number of iterations of a given algorithm, or its running time. Ideally, given a collection of actions  $\mathbf{x}$  played by players along the (approximate) equilibrium path computed, we wish to find the largest utility gain any player can have by deviating from this path, which we denote by  $\epsilon(\mathbf{x})$ . However, this computation is impossible in our setting, as it would need to consider all possible histories as well, whereas our approach and alternatives only return  $\mathbf{x}$  along the path of play (moreover, considering all possible histories is itself intractable).

We use the *global SPE regret* to evaluate our methods. It considers for each player  $i$  in level  $l$  a discrete grid of alternative actions, and uses best response dynamics to compute an approximate SPE of the level- $(l + 1)$  subgame to evaluate player  $i$ 's utility for each such deviation. This approach then returns the highest regret among all players computed in this way.

#### 4.1 DECENTRALIZED EPIDEMIC POLICY GAME

First, we consider DBI for solving a class of games inspired by hierarchical decentralized policy-making in the context of epidemics such as COVID-19 (Jia et al., 2021). The hierarchy has levels corresponding to the (single) federal government, multiple states, and county administrations under each state. Each player's action (policy) is a scalar in  $[0, 1]$  that represents, for example, the extent of social distancing recommended or mandated by a player (e.g., a state) for its administrative subordinates (e.g., counties). Crucially, these subordinates have considerable autonomy about setting their own policies, but incur a non-compliance cost for significantly deviating from recommendations made by the level immediately above (of course, non-compliance costs are not relevant for the root player). The full cost function of each player additionally includes infection prevalence within the geographic territory of interest to the associated entity (e.g., within the state), as well as the socio-economic cost of the policy itself (complete details are provided in Appendix C).

Since the actions are in a one-dimensional compact space and the depth of the hierarchy is at most 3, our baseline is the best response dynamics (BRD) algorithm proposed by Jia et al. (2021) (detailed in Appendix C), and we use *global regret* as a measure of efficacy in comparing the proposed DBI algorithm with BRD. The results of this comparison are shown in Figure 2 for two-level (government and states) and three-level (government, states, counties) variants of this game. We consider two-level games with 20 and 50 leaves (states), and three-level games with 2 players in level 2 (states) and 4 and 10 leaves (counties).




Figure 2: Global regret (top) and running time (bottom) for the decentralized epidemic policy game. Left (resp., right) column corresponds to result for games with 2(resp., 3) levels.

As we can see from the top plots in Figure 2, BRD can have poor convergence behavior in terms of global regret, whereas DBI appears to converge quite reliably to a path of play with a considerably lower global regret. Notably, the improvement in solution quality becomes more substantial as we increase the game complexity either in terms of scale (number of leaves) or in terms of the level of hierarchy (moving from 2- to 3-level games).

Running time (in seconds) demonstrates the relative efficacy of DBI even further (bottom plots in Figure 2). In particular, observe the significant increase in the running time of BRD as we increase the number of leaves. In contrast, DBI is far more scalable: indeed, even more than doubling the number of players appears to have little impact on its running time. Moreover, BRD is several orders of magnitude slower than DBI for the more complex games.

## 4.2 HIERARCHICAL PUBLIC GOODS GAMES

Next, we consider *hierarchical public goods games*. A conventional networked public goods game endows each player  $i$  with a utility function  $u_i(x_i, x_{-i}) = a_i + b_i x_i + \sum_j g_{ji} x_i x_j - c_i(x_i)$ , where  $g_{ji}$  is the impact of player  $j$  on player  $i$  (often represented as a weighted edge on a network), and  $x_i \in [0, 1]$  the level of investment in the public good by player  $i$  (Bramoullé & Kranton, 2007). We construct a 3-level hierarchical variant of such games by starting with the karate club network (Zachary, 1977) which represents friendships among 34 individuals. Level-2 nodes are obtained by partitioning the network into two (sub)clubs, with leaves (level-3 nodes) representing all the individuals. The utility of level-2 nodes is the sum of utilities of individual members of associated clubs, with the utility of the root being the sum of the utilities of all individuals. Furthermore, we introduce non-compliance costs with investment policies in the level immediately above, as we did in the decentralized epidemic policy game (Section 4.1). Further details are provided in Appendix C.2.




Figure 3: (a) Performance ( $\epsilon$ ) in the Public Goods Game (Section 4.2); the scatter points show the results of BRD with discretized factors 0.5, 0.2, 0.1, 0.05, and best response rounds 2, 3. (b), (c) Results on (1, 3, 6) hierarchical security games (Section 4.3); legend is shared.

Figure 3(a) presents the global regret as a function of running time for DBI (black line) and BRD with different levels of discretization (dots). We observe that DBI yields considerably lower regret in these games than BRD even as we discretize the latter finely. Moreover, DBI reaches smaller regret orders of magnitude faster than BRD.

## 4.3 HIERARCHICAL SECURITY GAMES

In the final set of experiments, we evaluate DBI on a hierarchical extension of *interdependent security games* (Bachrach et al., 2013). In these games,  $n$  defenders can each invest  $x_i \geq 0$  in security. If defender  $i$  is attacked, the probability that the attack succeeds is  $1/(1 + x_i)$ . Furthermore, defenders are interdependent, so that a successful attack on defender  $i$  cascades to defender  $j$  with probability  $q_{ji}$ . In the variant we adopt, the attacker strategy is a uniform distribution over defenders (e.g., the “attacker” is just nature, with attacks representing stochastic exogenous failures). The utility of the defender is the probability of surviving the attack less the cost of security investment.

We extend this simultaneous-move game to a hierarchical structure consisting of one root player (e.g., government), three level-2 players (e.g., sectors), and six leaf players (e.g., organizations). The policy-makers in the first two levels of the game recommend an investment policy to the level below, and aim to maximize total welfare (sum of utilities) among the leaf players in their subtrees. Just as in both hierarchical epidemic and public goods games, whenever a player in level  $l$  does not act according to the recommendation of their parent in level  $l - 1$ , they incur a non-compliance cost. Complete model details are deferred to Appendix C.3. We conduct experiments with two weights  $\kappa$  that determine the relative importance of non-compliance costs in the decisions of non-root players in the game:  $\kappa \in \{0.1, 0.5\}$ .

Figures 3(b) and 3(c) present the results of comparing DBI with BRD on this class of games, where BRD is again evaluated with different levels of action space discretization (note, moreover, that in this setting discretizing actions is not enough, since these are unbounded, and we also had to impose an upper bound). We can observe that for either value of  $\kappa$ , DBI yields high-quality SPE approximation (in terms of global SPE regret) far more quickly than BRD. In particular, when we use relatively coarse discretization, BRD is approximately an order of magnitude slower, and yields



significantly higher regret. In contrast, if we use finer discretization for BRD, global regret for BRD and DBI becomes comparable, but now BRD is several orders of magnitude slower. For example, DBI converges within several seconds, whereas if we discretize  $x_i$  into multiples of 0.02, BRD takes nearly 2 hours, while discretization at the level of 0.01 results in BRD taking nearly 7 hours.

## 5 CONCLUSION

We introduced a novel class of hierarchical games, proposed a new game-theoretic solution concept and designed an algorithm to compute it. We assume a specific form of utility dependency between players and our solution concept only guarantees local stability. Improvement on each of these two fronts is an interesting direction for future work.

Given the generality of our framework, our approach can be used for many applications characterized by a hierarchy of strategic agents e.g., pandemic policy making. However, our modeling requires the full knowledge of the true utility functions of all players and our analysis assumes full rationality for all the players. Although the model we have addressed here is already challenging, these assumptions are unlikely to hold in many real-world applications. Therefore, further analysis is necessary to fully gauge the robustness of our approach before deployment.

## ACKNOWLEDGMENTS

This work was supported in part by the US Army Research Office under MURI grant # W911NF-18-1-0208.

## REFERENCES

- Kareem Amin, Satinder Singh, and Michael Wellman. Gradient methods for Stackelberg security games. In *Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pp. 2–11, 2016.
- Yoram Bachrach, Moez Draief, and Sanjeev Goyal. Contagion and observability in security domains. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 1364–1371. IEEE, 2013.
- David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of  $n$ -player differentiable games. In *Thirty-Fifth International Conference on Machine Learning*, pp. 354–363, 2018.
- Nicola Basilico, Stefano Coniglio, and Nicola Gatti. Methods for finding leader-follower equilibria with multiple followers. In *Fifteenth International Conference on Autonomous Agents and Multiagent Systems*, pp. 1363–1364, 2016.
- Yann Bramoullé and Rachel Kranton. Public goods in networks. *Journal of Economic theory*, 135(1): 478–494, 2007.
- Benjamin Chasnov, Lillian Ratliff, Eric Mazumdar, and Samuel Burden. Convergence analysis of gradient-based learning in continuous games. In *Thirty-Sixth Conference on Uncertainty in Artificial Intelligence*, pp. 935–944, 2020.
- Stefano Coniglio, Nicola Gatti, and Alberto Marchesi. Computing a pessimistic stackelberg equilibrium with multiple followers: The mixed-pure case. *Algorithmica*, 82(5):1189–1238, 2020.
- Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Thirty-Second International Conference on Neural Information Processing Systems*, pp. 9256–9266, 2018.
- Asen Dontchev and Tyrrell Rockafellar. *Implicit functions and solution mappings*. Springer, 2009.
- Tanner Fiez, Benjamin Chasnov, and Lillian Ratliff. Implicit learning dynamics in Stackelberg games: Equilibria characterization, convergence analysis, and empirical study. In *Thirty-Seventh International Conference on Machine Learning*, pp. 3133–3144, 2020.

- Drew Fudenberg, Fudenberg Drew, David K Levine, and David K Levine. *The theory of learning in games*, volume 2. MIT press, 1998.
- Oded Galor. *Discrete dynamical systems*. Springer Science & Business Media, 2007.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *Twenty-Seventh Conference on Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Michael Hill and Frederic Varone. *The public policy process*. Routledge, 2021.
- Roger Horn and Charles Johnson. *Matrix analysis*. Cambridge University Press, 2012.
- Adam Ibrahim, Waiss Azizian, Gauthier Gidel, and Ioannis Mitliagkas. Linear lower bounds and conditioning of differentiable games. In *Thirty-Seventh International Conference on Machine Learning*, pp. 4583–4593, 2020.
- Feiran Jia, Aditya Mate, Zun Li, Shahin Jabbari, Mithun Chakraborty, Milind Tambe, Michael Wellman, and Yevgeniy Vorobeychik. A game-theoretic approach for hierarchical policy-making. *CoRR*, abs/2102.10646, 2021.
- Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *Thirty-Seventh International Conference on Machine Learning*, pp. 4880–4889, 2020.
- Michael Kearns, Michael Littman, and Satinder Singh. Graphical models for game theory. In *Seventeenth Conference on Uncertainty in Artificial Intelligence*, pp. 253–260, 2001.
- John Kelley. *General topology*. Springer, 1955.
- Alistair Letcher. On the impossibility of global convergence in multi-loss optimization. In *Ninth International Conference on Learning Representations*, 2021.
- Sven Leyffer and Todd Munson. Solving multi-leader-follower games. *Preprint ANL/MCS-P1243-0405*, 2005.
- Shu Li and Tamer Başar. Distributed algorithms for the computation of noncooperative equilibria. *Automatica*, 23(4):523–533, 1987.
- Tianyi Lin, Zhengyuan Zhou, Panayotis Mertikopoulos, and Michael Jordan. Finite-time last-iterate convergence for multi-agent learning in games. In *Thirty-Seventh International Conference on Machine Learning*, pp. 6161–6171, 2020.
- Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- Eric Mazumdar, Lillian Ratliff, and Shankar Sastry. On gradient-based learning in continuous games. *SIAM Journal on Mathematics of Data Science*, 2(1):103–131, 2020.
- Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1):465–507, 2019.
- Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of GANs. In *Thirty-First International Conference on Neural Information Processing Systems*, pp. 1823–1833, 2017.
- Vaishnavh Nagarajan and Zico Kolter. Gradient descent GAN optimization is locally stable. In *Thirty-First International Conference on Neural Information Processing Systems*, pp. 5591–5600, 2017.
- Thanh Nguyen, Arunesh Sinha, and He He. Partial adversarial behavior deception in security games. In *Twenty-Ninth International Joint Conference on Artificial Intelligence*, pp. 283–289, 2020.
- Martin J Osborne and Ariel Rubinstein. *An introduction to game theory*. Oxford university press New York, 2004.

- R Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55, 1975.
- Amirreza Shaban, Ching-An Cheng, Nathan Hatch, and Byron Boots. Truncated back-propagation for bilevel optimization. In *Twenty-Second International Conference on Artificial Intelligence and Statistics*, pp. 1723–1732, 2019.
- Michael Shub. *Global stability of dynamical systems*. Springer, 1987.
- Heinrich Von Stackelberg. *The theory of the market economy*. Oxford University Press, 1952.
- Kai Wang, Lily Xu, Andrew Perrault, Michael Reiter, and Milind Tambe. Coordinating followers to reach better equilibria: End-to-end gradient descent for stackelberg games. *CoRR*, abs/2106.03278, 2021.
- Yuanhao Wang, Guodong Zhang, and Jimmy Ba. On solving minimax optimization locally: A follow-the-ridge approach. In *Eighth International Conference on Learning Representations*, 2020.
- Stephen Wiggins. *Introduction to applied nonlinear dynamical systems and chaos*. Springer Science & Business Media, 2003.
- Bryan Wilder, Marie Charpignon, Jackson Killian, Han-Ching Ou, Aditya Mate, Shahin Jabbari, Andrew Perrault, Angel Desai, Milind Tambe, and Maimuna Majumder. Modeling between-population variation in COVID-19 dynamics in Hubei, Lombardy, and New York City. *Proceedings of the National Academy of Sciences*, 117(41):25904–25910, 2020.
- Wayne W Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, 33(4):452–473, 1977.

## A OMITTED DETAILS FROM SECTION 2

**Proposition 1.** *The SPE notion we defined for an SHG model corresponds exactly to the SPE of its extensive-form game (EFG) representation.*

*Proof.* First we show how to construct the EFG representation of an SHG. Since within level  $l$  players act simultaneously, we can designate a canonical player order for player in level  $l$ , say we label them as  $(l, 1), (l, 2), \dots, (l, n_l)$ . Then in the EFG representation, we sequentialize the playing order as always be  $(1, 1), (2, 1), (2, 2), \dots, (3, 1), \dots, (L, 1), \dots, (L, n_L)$ . For every  $(l, 1), l = 1, \dots, L$ , its information sets are all in the form of some singleton  $\{h_{<l}\}$  for some history  $h_{<l}$ . While for  $(l, n_l), l > 1$ , its information sets are sets that consist of some  $h_{<l}$  concatenated with all possible moves of  $(l, 1), \dots, (l, n_{l-1})$ . Then according to the classical definition of the subgame notion Selten (1975); Mas-Colell et al. (1995), the only well-defined subgames are those rooted at  $(l, 1), \forall l$  since these are the only ones with singleton information set. So every level- $l$  subgames corresponds to a subgame in the EFG representation, and every subgame in the EFG representation maps to a level- $l$  subgame of the EFG by investigating the player index. Then the SPE for both representations are also equivalent.  $\square$

## B CONVERGENCE ANALYSIS AND PROOFS

### B.1 CONVERGENCE ANALYSIS

Our task is to provide sufficient conditions for the DBI algorithm to converge to a stable point. A few notes are in order. First, as we remarked earlier, stable points of DBI are not guaranteed to be SPE just as stable points of gradient ascent are not guaranteed to be globally optimal with general non-convex objective functions. Furthermore, DBI algorithm entails what are effectively iterative better-response updates by players, and it is well-known that best response dynamic processes in games will in general lead to cycles (Fudenberg et al., 1998).

To begin, we observe that the gradient updates in DBI can be interpreted as a discrete dynamical system,  $\mathbf{x}^{t+1} = F(\mathbf{x}^t)$ , with  $F(\mathbf{x}^t) = (\mathbf{I} + \alpha G)(\mathbf{x}^t)$  where  $G$  is an update gradient vector. This discrete system can be viewed as an approximation of a continuous limit dynamical system  $\dot{\mathbf{x}} = G(\mathbf{x})$  (i.e., letting  $\alpha \rightarrow 0$ ). A standard solution concept for such dynamical systems is a *locally asymptotic stable point (LASP)*.

**Definition 1** (Galor (2007)). *A continuous (or discrete) dynamical system  $\dot{\mathbf{x}} = G(\mathbf{x})$  (or  $\mathbf{x}^{t+1} = F(\mathbf{x}^t)$ ) has a locally asymptotic stable point (LASP)  $\mathbf{x}^*$  if  $\exists \epsilon > 0, \lim_{t \rightarrow \infty} \mathbf{x}^t = \mathbf{x}^*, \forall \mathbf{x}^0 \in \mathbb{B}_\epsilon(\mathbf{x}^*)$ .*

There are well-known necessary and sufficient conditions for the existence of an LASP.

**Proposition 2** (Characterization of LASP (Wiggins, 2003, Theorem 1.2.5, Theorem 3.2.1)). *A point  $\mathbf{x}^*$  is an LASP for the continuous dynamical system  $\dot{\mathbf{x}} = G(\mathbf{x})$  if  $G(\mathbf{x}^*) = 0$  and all eigenvalues of Jacobian matrix  $\nabla_{\mathbf{x}} G$  at  $\mathbf{x}^*$  have negative real parts. Furthermore, for any  $\mathbf{x}^*$  such that  $G(\mathbf{x}^*) = 0$ , if  $\nabla_{\mathbf{x}} G$  has eigenvalues with positive real parts at  $\mathbf{x}^*$ , then  $\mathbf{x}^*$  cannot be an LASP.*

Note that an LASP of DBI is an action profile of all players that satisfies the first-order conditions, i.e., it has the property that no player can improve their utility through a local gradient update. While the existence of an LASP depends on game structure, we show that under Assumption 1, and as long as the sufficient conditions for LASP existence in Proposition 2 are satisfied, DBI converges to LASP. We defer all the omitted proofs to Appendix C.

**Proposition 3.** *Let  $\lambda_1, \dots, \lambda_d$  denote the eigenvalues of the updating Jacobian  $\nabla_{\mathbf{x}} G$  at an LASP  $\mathbf{x}^*$  and define  $\lambda^* = \arg \max_{i \in [d]} \text{Re}(\lambda_i) / |\lambda_i|^2$ , where  $\text{Re}$  is the real part operator. Then with a learning rate  $\alpha < -2\text{Re}(\lambda^*) / |\lambda^*|^2$ , and an initial point  $\mathbf{x}^0 \in \mathbb{B}_\epsilon(\mathbf{x}^*)$  for some  $\epsilon > 0$  around  $\mathbf{x}^*$ , DBI converges to an LASP. Specifically, if the choice of learning rate equals  $\alpha^*$  and the modulus of matrix  $\rho(\mathbf{I} + \alpha^* \nabla_{\mathbf{x}} G) = 1 - \kappa < 1$ , then the dynamics converge to  $\mathbf{x}^*$  with the rate of  $O((1 - \kappa/2)^t)$ .*

Proposition 3 states that there exists a region such that, if the initial point is in that region, then DBI will converge to an LASP. We next show that if we assume first-order Lipschitzness for the update rule, then we can also characterize the region of initial points which converge to an LASP.

**Proposition 4.** *Suppose  $G$  is  $L$ -Lipschitz.<sup>4</sup> Then for all  $\mathbf{x}^0 \in \mathbb{B}_{\kappa/2L}(\mathbf{x}^*)$ ,  $\epsilon > 0$  and after  $T$  rounds of gradient update, DBI will output a point  $\mathbf{x}^T \in \mathbb{B}_\epsilon(\mathbf{x}^*)$  as long as  $T \geq \lceil \frac{2}{\kappa} \log \|\mathbf{x}^0 - \mathbf{x}^*\| / \epsilon \rceil$  where  $\kappa$  is as defined in Proposition 3.*

We further show that through random initialization, the probability of reaching a *saddle point* is 0, which means that with probability 1, DBI converges to an LASP in which players are playing *local* best responses.

**Proposition 5.** *Suppose  $G$  is  $L$ -Lipschitz. Let  $\alpha < 1/L$  and define the saddle points of the dynamics  $G$  as  $\mathcal{X}_{sad}^* = \{\mathbf{x}^* \in \mathcal{X} \mid \mathbf{x}^* = (\mathbf{I} + \alpha G)(\mathbf{x}^*), \rho((\mathbf{I} + \alpha \nabla_{\mathbf{x}} G)(\mathbf{x}^*)) > 1\}$ . Also let  $\mathcal{X}_{sad}^0 = \{\mathbf{x}^0 \in \mathcal{X} \mid \lim_{t \rightarrow \infty} (\mathbf{I} + \alpha G)^t(\mathbf{x}^0) \in \mathcal{X}_{sad}^*\}$  denote the set of initial points that converge to a saddle point. Then  $\mu(\mathcal{X}_{sad}^0) = 0$ , where  $\mu$  is Lebesgue measure.*

While our convergence analysis does not guarantee convergence to an approximate SPE, our experiments show that DBI is in fact quite effective in doing so in practice.

*Proof of Proposition 3.* The learning dynamics of DBI can be written as  $\mathbf{x}^t = (\mathbf{I} + \alpha G)(\mathbf{x}^{t-1})$ . Since  $G$  has eigenvalues  $\lambda_1, \dots, \lambda_d$  at  $\mathbf{x}^*$ , the matrix  $\mathbf{I} + \alpha \nabla_{\mathbf{x}} G$  at a stationary point  $\mathbf{x}^*$  has eigenvalues  $1 + \alpha \lambda_1, \dots, 1 + \alpha \lambda_d$ . Since the set of LASPs is non-empty,  $Re(\lambda_i) < 0$  for all  $\lambda_i$ . Then for the choice of  $\alpha^*$  in the Proposition, the modulus of the Jacobian

$$\rho(\mathbf{I} + \alpha^* \nabla_{\mathbf{x}} G) = \max_{i \in [d]} |1 + \alpha^* \lambda_i| = \max_{i \in [d]} \sqrt{1 + 2\alpha^* Re(\lambda_i) + (\alpha^*)^2 |\lambda_i|^2} < 1.$$

Proofs for the convergence rate of  $O((1 - \kappa/2)^t)$  can be found in previous work (e.g., Fiez et al. (2020, Proposition F.1) or Wang et al. (2020, Proposition 4)). For completeness, we provide a proof. Since  $\rho((\mathbf{I} + \alpha \nabla_{\mathbf{x}} G)(\mathbf{x}^*)) = 1 - \kappa$ , according to Horn & Johnson (2012, Lemma 5.6.10), there exists a matrix norm  $\|\cdot\|$  such that  $\|\mathbf{I} + \alpha \nabla_{\mathbf{x}} G\| < 1 - \kappa + \epsilon$ , for  $\forall \epsilon > 0$ . We choose  $\epsilon = \frac{\kappa}{4}$ . The Taylor expansion of  $(\mathbf{I} + \alpha G)$  at  $\mathbf{x}^*$  is

$$(\mathbf{I} + \alpha G)(\mathbf{x}) = (\mathbf{I} + \alpha G)(\mathbf{x}^*) + (\mathbf{I} + \alpha \nabla_{\mathbf{x}} G)(\mathbf{x}^*)(\mathbf{x} - \mathbf{x}^*) + R(\mathbf{x} - \mathbf{x}^*),$$

where  $R(\mathbf{x} - \mathbf{x}^*) = o(\|\mathbf{x} - \mathbf{x}^*\|)$ . Let  $R_1(\mathbf{x} - \mathbf{x}^*) = \frac{1}{\alpha} R(\mathbf{x} - \mathbf{x}^*)$ . Then we have  $\lim_{\mathbf{x} \rightarrow \mathbf{x}^*} \frac{R_1(\mathbf{x} - \mathbf{x}^*)}{\|\mathbf{x} - \mathbf{x}^*\|} = 0$ . Then we can choose  $\delta > 0$  such that  $\|R_1(\mathbf{x} - \mathbf{x}^*)\| \leq \frac{\kappa}{4} \|\mathbf{x} - \mathbf{x}^*\|$  when  $\|\mathbf{x} - \mathbf{x}^*\| < \delta$ .

$$\begin{aligned} \|G(\mathbf{x}) - G(\mathbf{x}^*)\| &\leq \|\nabla_{\mathbf{x}} G(\mathbf{x}^*)(\mathbf{x} - \mathbf{x}^*)\| + \|R_1(\mathbf{x} - \mathbf{x}^*)\| \\ &\leq \|\nabla_{\mathbf{x}} G(\mathbf{x}^*)\| \|\mathbf{x} - \mathbf{x}^*\| + \frac{\kappa}{4} \|\mathbf{x} - \mathbf{x}^*\| \\ &\leq (1 - \frac{\kappa}{2}) \|\mathbf{x} - \mathbf{x}^*\|. \end{aligned}$$

This shows the operator  $(\mathbf{I} + \alpha G)$  is a contract mapping with contraction constant  $(1 - \frac{\kappa}{2})$ . Therefore the convergence rate is  $O((1 - \kappa/2)^t)$   $\square$

Before we present the proof of Proposition 4, we state the following lemma.

**Lemma 1.** *The update gradient vector  $G$  is  $L$ -Lipchitz if and only if  $\|\nabla_{\mathbf{x}} G\| \leq L$  at all  $\mathbf{x} \in \mathcal{X}$ .*

*Proof.* First we prove for "if" direction. Consider  $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$ ,

$$\begin{aligned} G(\mathbf{x}_2) - G(\mathbf{x}_1) &= \int_0^1 \nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau(\mathbf{x}_2 - \mathbf{x}_1))(\mathbf{x}_2 - \mathbf{x}_1) d\tau \\ \Rightarrow \|G(\mathbf{x}_2) - G(\mathbf{x}_1)\| &= \left\| \int_0^1 \nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau(\mathbf{x}_2 - \mathbf{x}_1))(\mathbf{x}_2 - \mathbf{x}_1) d\tau \right\| \\ &\leq \left\| \int_0^1 \nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau(\mathbf{x}_2 - \mathbf{x}_1)) d\tau \right\| \|\mathbf{x}_2 - \mathbf{x}_1\| \end{aligned}$$

<sup>4</sup>Formally, this means that  $\exists L > 0$  such that  $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}, \|G(\mathbf{x}) - G(\mathbf{x}')\|_2 \leq L \|\mathbf{x} - \mathbf{x}'\|_2$ .

$$\begin{aligned} &\leq \int_0^1 \|\nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau(\mathbf{x}_2 - \mathbf{x}_1))\| d\tau \|\mathbf{x}_2 - \mathbf{x}_1\| \\ &\leq L \|\mathbf{x}_2 - \mathbf{x}_1\|. \end{aligned}$$

Then we prove for the "only if" direction. Take  $\epsilon > 0$ , then for  $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$ ,

$$\begin{aligned} &\left\| \int_0^\epsilon (\nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau \mathbf{x}_2)) \cdot \mathbf{x}_2 d\tau \right\| = \|G(\mathbf{x}_1 + \epsilon \mathbf{x}_2) - G(\mathbf{x}_1)\| \leq \epsilon L \|\mathbf{x}_2\| \\ \Rightarrow \lim_{\epsilon \rightarrow 0} \frac{\left\| \int_0^\epsilon (\nabla_{\mathbf{x}} G(\mathbf{x}_1 + \tau \mathbf{x}_2)) \cdot \mathbf{x}_2 d\tau \right\|}{\epsilon \|\mathbf{x}_2\|} &= \frac{\|\nabla_{\mathbf{x}} G(\mathbf{x}_1) \cdot \mathbf{x}_2\|}{\|\mathbf{x}_2\|} \leq L. \end{aligned}$$

Since it holds for any  $\mathbf{x}_2$ , it must be  $\|\nabla_{\mathbf{x}} G(\mathbf{x}_1)\| \leq L$ . And since it applies for any  $\mathbf{x}_1 \in \mathcal{X}$ , this completes the proof.  $\square$

*Proof of Proposition 4.* From the proof from Proposition 3 notice that in order to make  $\|\mathbf{x}^t - \mathbf{x}^*\| \leq \epsilon$  we only have to let  $t \geq \lceil \frac{2}{\kappa} \log \|\mathbf{x}^0 - \mathbf{x}^*\| / \epsilon \rceil$  since now  $\|\mathbf{x}^t - \mathbf{x}^*\| \leq (1 - \frac{\kappa}{2})^t \|\mathbf{x}^0 - \mathbf{x}^*\| \leq \exp(-\kappa/2t) \leq \epsilon$ . Now we need to characterize the region of initial points that can converge to  $\mathbf{x}^*$  by characterizing the maximum possible radius of such initial point to  $\mathbf{x}^*$ . Recall in the proof of Proposition 3, this is captured by the parameter  $\delta$ . On one hand by using the Lipschitzness, we can bound the residual function by

$$\|R_1(\mathbf{x} - \mathbf{x}^*)\| \leq \int_0^1 \|\mathbf{I} + \alpha \nabla_{\mathbf{x}} G(\mathbf{x}^* + \tau(\mathbf{x} - \mathbf{x}^*)) - (\mathbf{I} + \alpha \nabla_{\mathbf{x}} G(\mathbf{x}^*))\| \|\mathbf{x} - \mathbf{x}^*\| d\tau \leq \frac{L}{2} \|\mathbf{x} - \mathbf{x}^*\|^2.$$

On the other hand to maintain this convergence rate we should let  $\|R_1(\mathbf{x} - \mathbf{x}^*)\| \leq \frac{\kappa}{4} \|\mathbf{x} - \mathbf{x}^*\|$ . Then we simply let  $\leq \frac{L}{2} \|\mathbf{x} - \mathbf{x}^*\|^2 \leq \frac{\kappa}{4} \|\mathbf{x} - \mathbf{x}^*\|$  we get an initial point should satisfy  $\|\mathbf{x} - \mathbf{x}^*\| \leq \frac{\kappa}{2L}$ .  $\square$

To prove Proposition 5, we need a few more machinery.

**Proposition 6.** *With  $\alpha < 1/L$ ,  $\mathbf{I} + \alpha G$  is a diffeomorphism.*

*Proof.* First we show  $(\mathbf{I} + \alpha G)$  is invertible. Suppose  $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$  such that  $\mathbf{x}_1 \neq \mathbf{x}_2$  and  $(\mathbf{I} + \alpha G)(\mathbf{x}_1) = (\mathbf{I} + \alpha G)(\mathbf{x}_2)$ . Then  $\mathbf{x}_1 - \mathbf{x}_2 = \alpha(G(\mathbf{x}_2) - G(\mathbf{x}_1))$ . And by  $\alpha < 1/L$  we have  $\|\mathbf{x}_1 - \mathbf{x}_2\| \leq \alpha L \|\mathbf{x}_1 - \mathbf{x}_2\| < \|\mathbf{x}_1 - \mathbf{x}_2\|$ , which is a contraction.

Next we show we show its invert function is well-defined on any point of  $\mathcal{X}$ . Notice that  $\rho(\alpha \nabla_{\mathbf{x}} G) \leq \|\alpha \nabla_{\mathbf{x}} G\| \leq \alpha L < 1$ . And notice that since the eigenvalues of  $\mathbf{I} + \alpha \nabla_{\mathbf{x}} G$  is the eigenvalues of  $\alpha \nabla_{\mathbf{x}} G$  plus 1, the only way that makes  $\det(\mathbf{I} + \alpha \nabla_{\mathbf{x}} G) = 0$  is to have one of the eigenvalues of  $\alpha \nabla_{\mathbf{x}} G$  to be  $-1$ . This contradicts to  $\rho(\alpha \nabla_{\mathbf{x}} G) < 1$ . Therefore by the implicit function theorem,  $(\mathbf{I} + \alpha G)$  is a local diffeomorphism on any point of  $\mathcal{X}$ , and therefore  $(\mathbf{I} + \alpha G)^{-1}$  is well defined on  $\mathcal{X}$ .  $\square$

**Theorem 2** (Center and Stable Manifold (Shub, 1987, Theorem III.7, Chapter 5)). *Suppose  $\mathbf{x}^* = h(\mathbf{x}^*)$  is a critical point for the  $C^r$  local diffeomorphism  $h : \mathcal{X} \rightarrow \mathcal{X}$ . Let  $\mathcal{X} = \mathcal{X}_s \oplus \mathcal{X}_u$ , where  $\mathcal{X}_s$  is the stable center eigenspace belonging to those eigenvalues of  $\nabla_{\mathbf{x}} h(\mathbf{x}^*)$  whose modulus is no greater than 1, and  $\mathcal{X}_u$  is the unstable eigenspace belonging to those whose modulus is greater than 1. Then there exists a  $C^r$  embedded disk  $W_{loc}^{cs}(\mathbf{x}^*)$  that is tangent to  $\mathcal{X}_s$  at  $\mathbf{x}^*$  called the local stable center manifold. Moreover there  $\exists \epsilon > 0$ ,  $h(W_{loc}^{cs}(\mathbf{x}^*)) \cap \mathbb{B}_\epsilon(\mathbf{x}^*) \subset W_{loc}^{cs}$  and  $\bigcap_{t=0}^\infty h^{-t}(\mathbb{B}_\epsilon(\mathbf{x}^*)) \subset W_{loc}^{cs}(\mathbf{x}^*)$ .*

*Proof of Proposition 5.* For  $\forall \mathbf{x}^* \in \mathcal{X}_{sad}^*$ , let  $\epsilon(\mathbf{x}^*) > 0$  be the radius of neighborhood provided by Theorem 2 for diffeomorphism  $h = \mathbf{I} + \alpha G$  and point  $\mathbf{x}^*$ . Then define  $\mathcal{B} = \bigcup_{\mathbf{x}^* \in \mathcal{X}_{sad}^*} \mathbb{B}_{\epsilon(\mathbf{x}^*)}(\mathbf{x}^*)$ . And since  $\mathcal{X}$  is a subset of Euclidean space so it is second-countable. By Lindelöf's Lemma Kelley (1955), which stated that every open cover there is a countable subcover, we can actually write  $\mathcal{B} = \bigcup_{i=1}^\infty \mathbb{B}_{\epsilon(\mathbf{x}_i^*)}(\mathbf{x}_i^*)$  for a countable family of saddle points  $\{\mathbf{x}_i^*\}_{i=1}^\infty \subseteq \mathcal{X}_{sad}^*$ . Therefore for  $\forall \mathbf{x}^0 \in \mathcal{X}_{sad}^*$  that converges to a saddle point, it must converge to  $\mathbf{x}_i^*$  for some  $i$ . And  $\exists T(\mathbf{x}^0) > 0$ , such that  $\forall t > T(\mathbf{x}^0)$ ,  $h^t(\mathbf{x}^0) \in \mathbb{B}_{\epsilon(\mathbf{x}_i^*)}(\mathbf{x}_i^*)$ . From Theorem 2 we have  $h^t(\mathbf{x}^0) \in W_{loc}^{sc}(\mathbf{x}_i^*)$ . Since  $h$  is a diffeomorphism on  $\mathcal{X}$  we have  $\mathbf{x}^0 \in h^{-t}(W_{loc}^{sc}(\mathbf{x}_i^*))$ . We furthermore union over all finite time

step  $\mathbf{x}^0 \in \bigcup_{t=0}^{\infty} h^{-t}(W_{loc}^{sc}(\mathbf{x}_i^*))$ . Then we have  $\mathcal{X}_{sad}^0 \subseteq \bigcup_{i=1}^{\infty} \bigcup_{t=0}^{\infty} h^{-t}(W_{loc}^{sc}(\mathbf{x}_i^*))$ . For each  $i$  since  $\mathbf{x}_i^*$  is a saddle point, it has an eigenvalue greater than 1, so the dimension of unstable eigenspace  $\dim(\mathcal{X}_u(\mathbf{x}_i^*)) \geq 1$ . Therefore the dimension of  $W_{loc}^{sc}(\mathbf{x}_i^*)$  is less than full dimension. This leads to  $\mu(W_{loc}^{sc}(\mathbf{x}_i^*)) = 0$  for any  $i$ . Again since  $h$  is a diffeomorphism,  $h^{-1}$  is locally Lipschitz which are null set preserving. Then  $\mu(h^{-t}(W_{loc}^{sc}(\mathbf{x}_i^*))) = 0$  for  $\forall i, t$ . And since the countable union of measure zero sets is measure zero,  $\mu(\bigcup_{i=1}^{\infty} \bigcup_{t=0}^{\infty} h^{-t}(W_{loc}^{sc}(\mathbf{x}_i^*))) = 0$ . So  $\mu(\mathcal{X}_{sad}^0) = 0$ .  $\square$

## C OMITTED DETAILS FROM SECTION 4

### C.1 DETAILS OF COVID-19 GAME EXPERIMENTS

**Subgame Perfect Equilibrium in an SHG** We here elaborate how we evaluate the regret of a profile. First let's imagine when solving a two-level games with  $n_2$  players at the second-level. Given  $(1, 1)$ 's action choice, the problem of computing SPE in the second-level is equivalent to computing a Nash equilibrium in the second-level. However, an exact pure Nash equilibrium may not exist. So which profile should we choose to propagate back to  $(1, 1)$ ? To resolve this issue, we define  $\varepsilon$ -Nash equilibrium.

**Definition 2.** For a simultaneous-move game, a profile  $\mathbf{x}^*$  is an  $\varepsilon$ -Nash if for any player  $n$ ,  $\forall \mathbf{x}'_n \in \mathcal{X}_n$ ,  $u_n(\mathbf{x}'_n, \mathbf{x}_{-n}^*) \leq u_n(\mathbf{x}^*) + \varepsilon$ .

In another word, an  $\varepsilon$ -Nash is a profile where for every player a unilateral deviation cannot offer benefit more than  $\varepsilon$  while fixing other's profile. In the context of our example for the two level game, given  $(1, 1)$ 's action, we select the profile with the minimum  $\varepsilon$  of the simultaneous-move game defined on level 2 as an SPE back to  $(1, 1)$ .

We now generalize this example to formally define the notion of  $\varepsilon$ -SPE in an SHG. First let us define  $\Phi_{l,i}(\mathbf{x}) \in \mathbb{R}^{d_L}$  that returns the equilibrated profile at level  $L$ . In this profile, leaves that are not descendants of  $(l, i)$  are fixed in  $\mathbf{x}$ , while  $LEAF(l, i)$  moved to a profile that corresponding to an SPE of  $\mathcal{G}_{l,i}(\mathbf{x})$  with an minimum  $\varepsilon$ . Then we define the  $\varepsilon_{l,i}(\mathbf{x})$  as  $\max_{\mathbf{x}'_{l,i} \in \mathcal{X}_{l,i}} u_{l,i}(\mathbf{x}'_{l,i}, \mathbf{x}_{PA(l,i)}, \Phi(\mathbf{x}'_{l,i})) - u_{l,i}(\mathbf{x}_{l,i}, \mathbf{x}_{PA(l,i)}, \Phi(\mathbf{x}_{l,i}))$  and define  $\varepsilon(\mathbf{x}) = \max_{l,i} \varepsilon_{l,i}(\mathbf{x})$  as the  $\varepsilon$  of profile  $\mathbf{x}$  in an SHG  $\mathcal{G}$ .

We next provide an approximate algorithm (see Algorithm 2) to compute  $\varepsilon$  in an SHG. Pay attention here that  $\mathbf{x} : \mathbf{x}_i \rightarrow \mathbf{x}'_i$  means replacing  $\mathbf{x}_i$  of  $\mathbf{x}$  with  $\mathbf{x}'_i$ . The functions in Algorithm 2 operate as follows. The procedure SEARCH takes a given joint profile  $\mathbf{x}$ , player index  $(l, i)$  and returns a best response profile for a given player  $(l, i)$ . Our approach is that for each action in  $\mathbf{x}'_{l,i} \in \mathcal{X}_{l,i}$ , we re-equilibrate subgames  $\mathcal{G}_{l,i}(\mathbf{x} : \mathbf{x}_{l,i} \rightarrow \mathbf{x}'_{l,i})$ , and then compute the corresponding payoff. In our actual implementation, we discretize  $\mathcal{X}_{l,i}$  in a bucket of grid points, and search within such bucket. The procedure RE-EQ returns the re-equilibrated profile of  $\mathcal{G}_{l,i}(\mathbf{x})$  given  $(l, i)$  and  $\mathbf{x}$ . The procedure SHG\_SOLVE solve a simultaneous-move game at level  $l + 1$ . It applies an iterative best response approach for CHD( $l, i$ ) to generate diverse profiles, and select the one with minimum possible  $\varepsilon$ . In each iteration, for each  $(l + 1, j) \in \text{CHD}(l, i)$ , it computes its best response action against the previous joint profile  $\mathbf{x}^{t-1}$ . And then in the next iteration it replace those descendant-profiles of  $(l + 1, j)$  in  $\mathbf{x}^{t-1}$  by computed re-equilibrated profiles when  $(l + 1, j)$  selected its best response. Then it will return the joint profile with the minimum  $\varepsilon$  found so-far.

To solve the whole game, we just call SHG\_SOLVE( $(0, 0), \mathbf{x}$ ), where  $\mathbf{x}$  is some other action profile. To compute the  $\varepsilon$  of a given profile, we just call COMPUTE\_  $\varepsilon$ , where it just compute the maximum unilateral deviation for every player using SEARCH to compute the best response action and payoff.

**Structured Game Model Inspired by COVID-19 Policy-Making** We will now describe in detail the particular subclass of SHGs that we studied in our experiments reported in Section 4.1. This class is based on the SHG proposed in Jia et al. (2021) which we describe in detail here. The exposition is in terms of a *cost function*  $C_i(\mathbf{x}_i, \mathbf{x}_{PA(i)}, \mathbf{x}_L)$  for each player, which is more natural in this context, rather than the payoff function  $u_{l,i}$  introduced in Section 2, with the understanding that  $u_i \equiv -C_i$ .

There are  $L = 3$  levels in the hierarchy such that player  $(1, 1)$  represents the federal government (or, simply, government), the players  $(2, i)$ ,  $i \in \{1, 2, \dots, n_2\}$  are state governments (or, simply, states),

---

**Algorithm 2** Procedures for computing  $\varepsilon$  of profile  $\mathbf{x}$  and an algorithm for computing an approximate SPE

---

**Input:** An SHG instance  $\mathcal{G}$

**Parameters:** Best response iterations  $T$

```

procedure SEARCH( $\mathbf{x}, (l, i)$ )
  for  $\mathbf{x}'_{l,i} \in \mathcal{X}_{l,i}$  do
     $u_{l,i}^{re-eq}(\mathbf{x}'_{l,i}), \varepsilon_{DES(l,i)}(\mathbf{x}'_{l,i}), \mathbf{x}^{re-eq(l,i)}(\mathbf{x}'_{l,i}) \leftarrow \text{RE-EQ}((l, i), \mathbf{x} : \mathbf{x}_{l,i} \rightarrow \mathbf{x}'_{l,i})$ 
  Find  $\mathbf{x}^*_{l,i} \leftarrow \arg \max_{\mathbf{x}'_{l,i}} u_{l,i}^{re-eq}(\mathbf{x}'_{l,i})$ 
  return  $\mathbf{x}^*_{l,i}, u_{l,i}^{re-eq}(\mathbf{x}^*_{l,i}), \mathbf{x}^{re-eq(l,i)}(\mathbf{x}^*_{l,i})$ 
end procedure
procedure RE-EQ( $\mathcal{G}, \mathbf{x}, (l, i)$ )
   $\varepsilon \leftarrow 0$ 
  if  $l < L$  then
     $\mathbf{x}, \varepsilon \leftarrow \text{SHG\_SOLVE}((l, i), \mathbf{x})$ 
  return  $u_{l,i}(\mathbf{x}), \varepsilon, \mathbf{x}$ 
end procedure
procedure SHG_SOLVE( $(l, i), \mathbf{x}$ )
   $\mathbf{x}^0 \leftarrow \mathbf{x}$ 
   $\forall (l+1, j) \in \text{CHD}(l, i)$ , replace its action in  $\mathbf{x}^0$  with other random initialization
  for  $t = 1, 2, \dots, T$  do ▷ level-wise best response
    for  $(l+1, j) \in \text{CHD}(l, i)$  do
       $u_{l+1,j}^{re-eq}, \varepsilon_{DES(l+1,j)}, \mathbf{x}^{re-eq(l+1,j)} \leftarrow \text{RE-EQ}(\mathcal{G}, \mathbf{x}^{t-1}, (l+1, j))$ 
       $\mathbf{x}'_{l+1,j}, u_{l+1,j}^{re-eq-max}, \mathbf{x}^{re-eq-max(l+1,j)} \leftarrow \text{SEARCH}(\mathbf{x}^{t-1}, (l+1, j))$ 
       $\varepsilon_{l+1,j}^{t-1} = \max\{\varepsilon_{DES(l+1,j)}, u_{l+1,j}^{re-eq-max} - u_{l+1,j}^{re-eq}\}$ 
     $\mathbf{x}^t \leftarrow \mathbf{x}^{t-1}$ 
    for  $\forall (l+1, j) \in \text{CHD}(l, i)$  do
      Replace dimensions of  $\mathbf{x}^t$  belonging to  $\text{DES}(l+1, j)$  with the ones in
       $\mathbf{x}^{re-eq-max(l+1,j)}$ 
       $\varepsilon_{l+1}^{t-1} \leftarrow \max_i \varepsilon_{l+1,j}^{t-1}$ 
     $t^* \leftarrow \arg \min_t \varepsilon_{l+1}^t$ 
  return  $\mathbf{x}^{t^*}, \varepsilon_{l+1}^{t^*}$ 
end procedure
procedure COMPUTE_ε( $\mathbf{x}$ )
  for  $(l, i)$  do
     $\mathbf{x}^*_{l,i}, u_{l,i}^*, \mathbf{x}^{re-eq(l,i)} \leftarrow \text{SEARCH}(\mathbf{x}, (l, i))$ 
     $\varepsilon_{l,i} \leftarrow u_{l,i}^* - u_{l,i}(\mathbf{x})$ 
  return  $\max_{l,i} \varepsilon_{l,i}$ 
end procedure

```

---

and the players  $(3, i)$ ,  $i \in \{1, 2, \dots, n_3\}$  are county governments (or, simply, counties) partitioned into groups such that each group shares a single state as a parent.

Each player  $(i)$  takes a bounded, scalar action  $x_i \in [0, 1]$  which is a *social-distancing factor* that (multiplicatively) reduces the proportion of post-intervention contacts among individuals — a lower number implies a stronger policy intervention, hence a lower number of infections but a higher cost of implementation (see below). The actions taken by counties represent policies that get *actually implemented* (hence directly impact the realized cost of every player — one of the defining attributes of SHGs) while those taken by the government and states are *recommendations*. Similar to Jia et al. (2021), we also study a restricted variant where each county is non-strategic and constrained to comply with the action (recommendation) of its parent-state, effectively reducing the model to a 2-level hierarchy. We call this special case a *two-level game* (Figure 2(a)) and the more general model a *three-level game* (Figure 2(b)).

The cost function of each player  $(i)$  has, in general, three components: a *policy impact cost*  $C_i^{\text{inc}}(\mathbf{x}_L)$  which we will elaborate on below; a *policy implementation cost*  $C_i^{\text{dec}}(\mathbf{x}_L)$ , e.g. economic and



psychological costs of a lockdown; and, for each player in levels  $l > 1$ , a *non-compliance cost*  $\mathbf{C}_i^{\text{NC}}(\mathbf{x}_i, \mathbf{x}_{\text{PA}(i)})$ , a penalty incurred by a policy-maker for deviating from the recommendation of its parent in the hierarchy (e.g., a fine, litigation costs, or reputation harm).

Let  $N_{(3,i)} > 0$  denote the fixed population under the jurisdiction of county  $(3, i)$  for every  $i \in \{1, 2, \dots, n_3\}$ . By construction, the population under each state is given  $N_{(2,i)} = \sum_{(3,j) \in \text{CHD}(2,i)} N_{(3,j)}$  and the population under the government is  $N_{(1,1)} = \sum_{i=1}^{n_2} N_{(2,i)}$ . We next define the expressions for each component of the cost function.

**Policy Impact Cost:** This cost component is a quadratic closed-form approximation to the agent-based model introduced by Wilder et al. (2020). This is inspired by the infection cost computation approach in Jia et al. (2021) but they used a different closed-form approximation. For each county  $(3, i)$ , let  $N_{(3,i)}^{\text{init}}$  denote the number of infected individuals within the population of the county prior to policy intervention; thus, the number of *post-intervention susceptible* individuals is  $(N_{(3,i)} - N_{(3,i)}^{\text{init}})x_{3,i}$ . Another parameter in the game is the *transport matrix*  $R = \{r_{aa'}\}_{a,a' \in (3,1),(3,2), \dots, (3,n_3)}$ , where  $r_{aa'} \geq 0$  is the proportion of the population of county  $a'$  that is active in county  $a$  in the absence of an intervention. Thus, in the number of post-intervention infected individuals of county  $a'$  that is active in county  $a$  is  $r_{aa'} N_{a'}^{\text{init}} x_{a'}$ . The last parameter in the model is  $M$ , the average number of contacts with active individuals that a susceptible individual makes, and finally  $\mu$  is the probability that a susceptible individual gets infected upon contact with an active infected individual is  $\mu \in (0, 1)$ .

Putting these together, the policy impact cost is defined by the fraction of post-intervention infected individuals in county  $a = (3, i)$ ,  $i \in \{1, 2, \dots, n_3\}$ :

$$\mathbf{C}_a^{\text{inc}}(\mathbf{x}_L) = \mu M x_a \frac{N_a - N_a^{\text{init}}}{N_a^2} \left( \sum_{a'} r_{aa'} N_{a'}^{\text{init}} x_{a'} \right).$$

For a higher-level player  $(i)$ ,

$$\mathbf{C}_i^{\text{inc}}(\mathbf{x}_L) = \frac{1}{N_{(i)}} \sum_{a \in \text{CHD}(i)} N_a \mathbf{C}_a^{\text{inc}}(\mathbf{x}_L).$$

**Policy Implementation Cost:** For each county  $(3, i)$ , the policy implementation cost is given by

$$\mathbf{C}_{3,i}^{\text{dec}}(\mathbf{x}_L) = 1 - x_{3,i}.$$

For a higher-level player  $(i)$ ,

$$\mathbf{C}_i^{\text{dec}}(\mathbf{x}_L) = \frac{1}{N_{(i)}} \sum_{a \in \text{CHD}(i)} N_a \mathbf{C}_a^{\text{dec}}(\mathbf{x}_L).$$

**Non-Compliance Cost:** The non-compliance cost of player  $(i)$  for  $l \in \{2, 3\}$  is given by Euclidean distance between its action and that of its parent:

$$\mathbf{C}_i^{\text{NC}}(x_i, x_{\text{PA}(i)}) = (x_i - x_{\text{PA}(i)})^2.$$

Finally, each player  $(i)$  for  $l > 1$  has an idiosyncratic set of *weights*  $\kappa_i \geq 0$  and  $\eta_i \geq 0$  that trade its three cost components off against each other via a convex combination, and account for differences in ideology; the overall cost of such a player is given by

$$\mathbf{C}_i(x_i, x_{\text{PA}(i)}, \mathbf{x}_L) = \kappa_i \mathbf{C}_i^{\text{inc}}(\mathbf{x}_L) + \eta_i \mathbf{C}_i^{\text{dec}}(\mathbf{x}_L) + (1 - \kappa_i - \eta_i) \mathbf{C}_i^{\text{NC}}(x_i, x_{\text{PA}(i)}). \quad (7)$$

The player  $(1, 1)$  obviously has no non-compliance issues, hence it has only one weight  $\kappa_{1,1} > 0$ , its overall cost being

$$\mathbf{C}_{1,1}(x_i, \mathbf{x}_L) = \kappa_{1,1} \mathbf{C}_{1,1}^{\text{inc}}(\mathbf{x}_L) + (1 - \kappa_{1,1}) \mathbf{C}_{1,1}^{\text{dec}}(\mathbf{x}_L).$$

In our experiments, we set  $r_{aa'} = 1/n_3$  for every pair of counties  $(a, a')$ ,  $M = 20$  and  $\mu = 0.3$ .

**Experimental Setup and Further Results** We formally define the algorithm BRD as the procedure  $\text{SHG\_SOLVE}((0, 0), x)$  for some randomly initialized  $x$ . We start by defining the concept of one full *algorithm iteration* for each of DBI and BRD. For DBI(1, 20), DBI(1, 50), DBI(1, 2, 4), DBI(1, 2, 10), one algorithm iteration consists of 50 steps of gradient ascent, with a learning rate of 0.01. For BRD(1, 20) and BRD(1, 50), one algorithm iteration consists of 20 iterations of level-wise best response during the recursive procedure in  $\text{SHG\_SOLVE}$ ; for BRD(1, 2, 4) and BRD(1, 2, 10), it corresponds to 500 and 200 iterations of level-wise best response, respectively. For DBI we adopt a projector operator that project the resulted action into the nearest point in  $[0, 1]$ .

We discretize each action space uniformly into 101 grid points for two-level games, and 11 grid points for three-level games. We let  $T = 100$  for BRD(1, 20) and BRD(1, 50) and  $T = 20$  for BRD(1, 2, 4) and BRD(1, 2, 10). In the three-level experiments, the  $\kappa$  is set to be 0.5 for counties and the states and 0.8 for the government. The  $\eta$  is set to be 0.2 for the states, 0.3 for the counties in (1,2,4) setting, and 0.2 for (1,2,10) experiment. In the two-level experiments, the  $\kappa$  is set to be 0.2 for the government, 0.5 for counties and the states. The  $\eta$  is set to be 0.2 for counties and states.

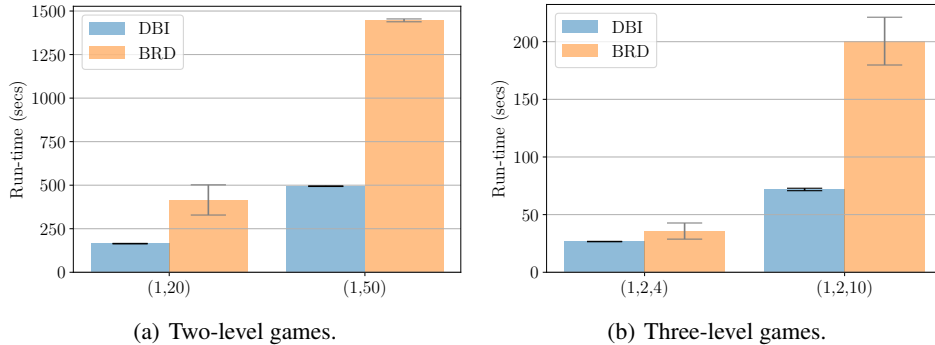


Figure 4: The run-time (in secs) on COVID-19 Game Models.  $T = 20$  for BRD experiments.

To further study the scalability of the BRD and DBI algorithms, we compare the run-time that (1) the DBI algorithm converges and (2) the BRD algorithm terminates. The convergence of DBI is defined as the convergence of the action profile. When DBI converges, the action profile remains unchanged. In the projected gradient descent method, if all the gradients go to zero, DBI converges. Besides, it is also likely to hit the boundary of the constraints. The DBI is possible to converge with a non-zero gradient norm. The BRD algorithm terminates either  $T$  achieves, or  $\epsilon$  goes to zero for all players. Under these conditions, the two algorithms find their optimal solutions. Figure 4 demonstrates the run-time results. We conduct each experiment four times with different random seeds. In two-level problems, the DBI algorithm is more than two times faster than BRD algorithms. Although the action spaces are discretized to 11 grid points for three-level games rather than 101 grids, DBI algorithms still perform better. In practice, discretizing the action spaces in 101 grids for three-level games is computationally intensive. The performance and run-time of the BRD algorithm are more dependent on randomization and the initial points for the best responses at each level. When we face many players or multiple levels, the DBI algorithm is a natural choice. The DBI algorithm is significantly more efficient and more stable than the BRD algorithm.

## C.2 DETAILS OF HIERARCHICAL PUBLIC GOODS GAME EXPERIMENTS

We extend the described public good game on Zachary’s Karate club network to a (1-2-34) hierarchical game by introducing the non-compliance costs similar to the COVID-19 game model. The overall utility of the club individuals should be the combination of public good utility  $u_i$  and  $C_i^{\text{NC}}(x_i, x_{\text{PA}(i)}) = (x_i - x_{\text{PA}(i)})^2$ .

$$U_{3,i}(x_i, x_{\text{PA}(i)}, \mathbf{x}_L) = (1 - \kappa_{3,i})u_i(x_i, \mathbf{x}_L) - \kappa_{3,i}C_i^{\text{NC}}(x_i, x_{\text{PA}(i)}). \quad (8)$$

In the second level, the administrator and instructor’s utility is composed of the total public good utility of their group members, and their own non-compliance costs.

$$U_{2,i}(x_i, x_{\text{PA}(i)}, \mathbf{x}_L) = (1 - \kappa_{2,i}) \sum_{j \in \text{CHD}(i)} u_j(x_j, \mathbf{x}_L) - \kappa_{2,i} C_i^{\text{NC}}(x_i, x_{\text{PA}(i)}). \quad (9)$$

The root player’s utility is considered as the social welfare of the whole club.

$$U_{1,1}(x_i, x_{\text{PA}(i)}, \mathbf{x}_L) = \sum_j u_j(x_j, \mathbf{x}_L) \quad (10)$$

In our experiments, we set  $\kappa_{2,i} = \kappa_{3,j} = 0.5$ ,  $\forall i, j$ . Other parameters of the public good utility are set to be  $a_i = 0, b_i = 1, c_i = 6$ ,  $\forall i$ . We use learning rate 0.1 in DBI. We project the results to  $[0, 1]$ . To compare with the DBI, we use BRD with discretized factors 0.5, 0.2, 0.1, 0.05, and best response rounds 2, 3.

### C.3 DETAILS OF HIERARCHICAL SECURITY GAME EXPERIMENTS

Let  $a = (a_1, \dots, a_n)$  be the attacker strategy, the probability distribution over which defenders are attacked. In our experiments,  $a$  is set to follow a logit distribution  $\text{softmax}(\lambda(1 - \mathbf{x}_L))$ , with defenders having lower security investment more likely to be attacked. Let  $c_i(x_i) = c_i x_i$  be the cost of security investment for defender  $i$ , which is assumed to be linear. The utility of defender  $i$  is then as follows:

$$u_i(x_i, \mathbf{x}_L) = a_i \frac{x_i}{1 + x_i} + \sum_{j \neq i} a_j \left(1 - q_{ji} \frac{1}{1 + x_j} \frac{1}{1 + x_i}\right) - c_i(x_i)$$

We then build the (1,3,6) structured hierarchical game by introducing the non-compliance costs  $C_i^{\text{NC}}(x_i, x_{\text{PA}(i)}) = (x_i - x_{\text{PA}(i)})^2$ . The utility of each agent in different levels are the same as that of public good game (Eq. 8, 9, 10).

In the experiments, we consider the agents who almost obey the parent’s suggestion ( $\kappa = 0.1$ ) and the somewhat selfish agents ( $\kappa = 0.5$ ). The influence probability  $q_{ji}$ ,  $\forall i, j$  are set to be equal among the agents, and other parameters are set to  $c_i = 0.2, \lambda = 5$ . We use a learning rate of 0.1 in DBI. To compare with the DBI, we use BRD with discretized factors 0.2, 0.1, 0.05 and best response rounds 2, 3. Since there is no upper bound of the strategy profile, we set  $x_i \leq 1.0$  according to the performance. Then we search the strategy space  $x \in [0, 1]$  in the BRD execution.