

Learning Game-theoretic Models from Aggregate Behavioral Data with Applications to Vaccination Rates in Public Health

Hau Chan

University of Nebraska-Lincoln
Lincoln, NE
hchan3@unl.edu

Luis E. Ortiz

University of Michigan-Dearborn
Dearborn, MI
leortiz@umich.edu

ABSTRACT

A survey is a common method to elicit behavioral data. The collected data provides a noisy representation of the actions of a sampled population. Direct access to individual responses is rare, for many obvious reasons. Instead, most of us would only have access to aggregated information about the percentage of individuals who reportedly took certain actions. Public-health data on populations' vaccination rates collected by government officials is such example and will be the focus of this paper. Naturally, behavioral data capture implicit interdependencies governing the decision-making process of the sampled population. In this work, we undertake the challenging task of uncovering such independencies of the data and use computational game theory (CGT) to model data as the result of distributed decision-making at the reported granularity level (e.g., nations, states, districts, and towns). Indeed, CGT has increasingly gained popularity as both a formal and practical framework in which to study the potential effect of policy making of agents in a variety of settings, like the vaccination setting we study here. To achieve our task, in this paper, we posit the view of aggregated behavioral data as jointly randomized, or mixed, strategies of multiple agents. We propose a novel general machine-learning approach to infer game-theoretic models from a potentially noisy dataset of mixed strategies. Our goal is to learn instantiations of game-theoretic models from the data that would best explain and compactly represent the global behavior of the population within a given hypothesis class of games. Ultimately, we want to employ the learned models for policy analysis on the underlying system as a whole, which cannot be achieved using other existing approaches. We illustrate our framework in a health-policy setting using publicly available data on vaccination rates in the continental USA.

ACM Reference Format:

Hau Chan and Luis E. Ortiz. 2018. Learning Game-theoretic Models from Aggregate Behavioral Data with Applications to Vaccination Rates in Public Health. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018*, IFAAMAS, 9 pages.

1 INTRODUCTION

Computational game theory (CGT) has become an important tool in science and engineering. Its use has increased considerably in recent decades. A major reason for its popularity comes from the strong mathematical foundation it provides for settings of distributed

decision-making. Security games (see, e.g., a survey by [21]), network security games (see, e.g., surveys by [26] and [23]), influence games [17], and vaccination games [15], among others, exemplify models used to describe and help predict joint global behavior of individual entities in complex systems. Largely motivated by datasets from the USA's *Center for Disease Control and Prevention (CDC)*, here we apply CGT, AI, and ML to a vaccination setting at a new scale and abstraction level.

The CDC collects and reports aggregate data about vaccination rates, along with standard deviations, for each state (of the U.S.) yearly. Each vaccination percentages represent the state-wide behavior of the people living in the State. Alternatively, we can view each state's vaccination percentage as a proxy measure of the state government's achievement from effort to raise its population vaccination rate for some disease or epidemic (e.g., influenza and Ebola). For instance, this can be done by increasing the awareness through state CDC newsletters and websites¹. As a result, we can view those vaccination rates of the states as the joint-behavior of the states (i.e., the outcome of their efforts). Given these state vaccination rates, we want to understand how the epidemic vaccination decisions of the states affect each other by modeling the strategic interaction as vaccination games (the variant used here is defined in Section 4).

More generally, our main interest is to infer interdependent characteristics inherent in the system from just the state vaccination data we have available. We seek to address a couple of fundamental scientific and engineering questions for this type of data. One is the "the learning question," how much can we learn in terms of compact representation of global behavior exclusively from the aggregate reports we have available about each entity in the system in isolation, by exploiting implicit interdependencies encoded in those pieces of aggregate information? The other is the "inference question," what type of questions we can study, particularly in terms of state policy making and evaluation? In this paper, we mostly focus on the learning question, and provide an illustration addressing the inference question.

Contribution. We conclude the introduction with a summary of our contributions. Our interest in this work is learning games from observed high-dimensional mixed-strategy data. In contrast to previous work, in which the data are the actions or pure strategies of the players [16], we are dealing with data that summarizes the actions of all the individuals within a state's population. In our model, we view each vaccination rate as representing the mixed-strategy of each state agent. In game-theoretic terms, we view these probabilities collectively as possibly approximate mixed-strategy Nash equilibrium (MSNE) to account for noises. In particular, we

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹patch.com/minnesota/saintpaul/flu-rising-state-health-officials-urge-flu-shots

- propose and introduce a machine learning (generative) framework to learn a game given the behavioral data;
- use our framework to derive a heuristic to learn α -IDS games given the CDC vaccination data; and
- experimentally show that our framework and learning heuristic are effective for inferring α -IDS games for understanding “global” strategic behavior and illustrate policy analysis of the states in the U.S.

Related Work. The closest work to ours is that of Honorio and Ortiz [16], in which they provide a general machine-learning framework to learn the structure and parameters of games from discrete behavioral data (e.g., “Yes/No”-type responses). Moreover, they demonstrate their framework on learning some classes of games. We refer the reader to the related-work section of Honorio and Ortiz [16] for a more detailed discussion. For brevity, all other previous methods assume that the actions and payoffs are observable in the data [6, 7, 9, 12, 28–30] while others are interested in predicting future behavior from the past behavior (system dynamics) [20, 31].

2 PRELIMINARIES

We refer the reader to [11] for an introduction to basic concepts in game theory. Let $V = \{1, 2, \dots, n\}$ be a *set of players/agents*. For each $i \in V$, let A_i be the *set of actions/pure-strategies* available to i , and $A = \times_{i \in V} A_i$ the *set of joint-actions/joint-pure-strategies*. Denote by $u_i : A \rightarrow \mathbb{R}$ the *payoff of i* for each joint-action in A . Similarly, let X_i be the *set of mixed-strategies of i* , and $X \equiv \times_{i \in V} X_i$ be the *set of joint-mixed-strategies*, which is a probability simplex over A_i . Given the context of this paper and the application domain (i.e., the CDC dataset), we assume that $A_i = \{0, 1\}$ for all i , so that we can represent $X_i = [0, 1]$, and interpret each $x_i \in X_i$ as *the probability that agent i plays $a_i = 1$* . We also assume that for all i , the maximum and minimum payoff value of u_i is 1 and 0, respectively. Using the standard abuse of notation, we denote by $u_i(x) \equiv \mathbf{E}_{a \sim x}[u_i(a)] = \sum_{a \in A} \left[\prod_{j \in V} x_j^{a_j} (1 - x_j)^{(1-a_j)} \right] u_i(a)$ is the *expected payoff of player i with respect to $x \in X$* . Given $\epsilon \geq 0$, a mixed-strategy profile $x^* \in X$ is an ϵ -approximate mixed-strategy Nash equilibrium, or ϵ -MSNE for short, of a non-cooperative game [27], defined by the tuple $(V, A, \{u_i\}_{i \in V})$ if, for each player $i \in V$, $u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*) - \epsilon$ for all $x_i \in X_i$. Moreover, we denote the set of all ϵ -MSNE of a game \mathcal{G} as $\mathcal{NE}^\epsilon(\mathcal{G}) \equiv \{x^* \in [0, 1]^n \mid \forall i, u_i(x_i^*, x_{-i}^*) \geq u_i(0, x_{-i}^*) - \epsilon \text{ and } u_i(x_i^*, x_{-i}^*) \geq u_i(1, x_{-i}^*) - \epsilon\}$. A 0-MSNE is an *exact MSNE*, which always exists for any non-cooperative game [24].

3 A FRAMEWORK TO LEARN GAMES

We now present our general framework to learn games from joint-mixed-strategy behavioral data.

3.1 A Generative Model for Behavioral Data

Motivated in part by the CDC data, we propose a generative model of behavioral data over the set of *mixed-strategy* profiles. We build on the work of [16], in which they provide a general machine-learning framework to learn the structure and parameters of games from *discrete* behavioral data (e.g., “Yes/No”-type responses). This

will present a number of challenges, which we will discuss after formally presenting our proposed generative model.

We assume (as in [16]) the statistical process generating the data is a simple mixture model: i.e., with probability $q \in (0, 1)$, the process generates/outputs a mixed-strategy profile x by drawing uniformly at random from the set $\mathcal{NE}^\epsilon(\mathcal{G})$; with probability $1 - q$, the process generates a mixed-strategy profile x by drawing uniformly at random from its complement set $[0, 1]^n - \mathcal{NE}^\epsilon(\mathcal{G})$, the complement of $\mathcal{NE}^\epsilon(\mathcal{G})$. Said differently, our generative model of behavioral data for mixed-strategy profiles is a mixture model with mixture parameter q and mixture components defined in terms of the approximation parameter $\epsilon > 0$ and a game \mathcal{G} . Even though we view each reported vaccination rate as a mixed-strategy, we do not assume that it corresponds to an exact MSNE. For instance, in the context of state-level vaccination rates, noise is common on the type of aggregate data that the CDC collects and reports. As a result, ϵ is a way to account for noises in the data.

Let μ be the standard *Borel measure*. (We refer the reader to a standard textbook [1] for an introduction to measure-theoretic concepts.) Note that, in our context, because the outcome space is $[0, 1]^n$ and $\mu([0, 1]^n) = 1$, we can view the Borel measure μ as the *uniform probability measure*. More formally, the *probability density function (PDF) f* for the generative model with parameters $(q, \mathcal{G}, \epsilon)$ over the hypercube of joint-mixed-strategies $[0, 1]^n$ is

$$f_{(q, \mathcal{G}, \epsilon)}(x) \equiv q \frac{\mathbb{1}[x \in \mathcal{NE}^\epsilon(\mathcal{G})]}{\mu(\mathcal{NE}^\epsilon(\mathcal{G}))} + (1 - q) \frac{\mathbb{1}[x \notin \mathcal{NE}^\epsilon(\mathcal{G})]}{1 - \mu(\mathcal{NE}^\epsilon(\mathcal{G}))}, \quad (1)$$

for all $x \in [0, 1]^n$. Our generative model was adapted from the early work in [16] and it has already been shown to be effective in learning games. Note that our definition of $f_{(q, \mathcal{G}, \epsilon)}$ requires us to deal with a few subtle but crucial potential measure-theoretic obstacles that [16] did not have to face because their generative model was a probability mass function. Key among those obstacles is that our model requires the (Borel) *measurability* of $\mathcal{NE}^\epsilon(\mathcal{G})$. Another potential obstacle is the *a priori* possibility that $\mathcal{NE}^\epsilon(\mathcal{G})$ might have measure zero. The following lemma gets us around the first roadblock.

LEMMA 3.1. *The set $\mathcal{NE}^\epsilon(\mathcal{G})$ is Borel μ -measurable for any game \mathcal{G} and any $\epsilon \geq 0$.*

PROOF. Recall that $\mathcal{NE}^\epsilon(\mathcal{G})$

$$\begin{aligned} &\equiv \{x^* \mid \forall i, u_i(x_i^*, x_{-i}^*) \geq u_i(0, x_{-i}^*) - \epsilon \\ &\quad \text{and } u_i(x_i^*, x_{-i}^*) \geq u_i(1, x_{-i}^*) - \epsilon\} \\ &\equiv \mathcal{BR}_1^\epsilon(\mathcal{G}) \cap \mathcal{BR}_2^\epsilon(\mathcal{G}) \cap \dots \cap \mathcal{BR}_n^\epsilon(\mathcal{G}), \end{aligned}$$

where

$$\begin{aligned} \mathcal{BR}_i^\epsilon(\mathcal{G}) &= \{x \mid u_i(x_i, x_{-i}) \geq u_i(0, x_{-i}) - \epsilon \\ &\quad \text{and } u_i(x_i, x_{-i}) \geq u_i(1, x_{-i}) - \epsilon\} \end{aligned}$$

for $i \in V$. The $\mathcal{BR}_i^\epsilon(\mathcal{G})$ of each player i is formed by two continuous multilinear/polynomial inequalities and $\mathcal{BR}_i^\epsilon(\mathcal{G}) \subseteq [0, 1]^n$. It follows that $\mathcal{NE}^\epsilon(\mathcal{G}) \subseteq [0, 1]^n$ is a semialgebraic set. The semialgebraic set is closed, and therefore $\mathcal{NE}^\epsilon(\mathcal{G})$ is measurable. \square

It turns out that we can also get around the second obstacle.

PROPOSITION 3.2. *For any $\epsilon > 0$, we have $\mu(\mathcal{NE}^\epsilon(\mathcal{G})) > 0$.*

While it is possible that the measure of $\mathcal{NE}^\epsilon(\mathcal{G})$ is zero for $\epsilon = 0$, its measure will be positive for $\epsilon > 0$. The key to show this fact is to realize that there is at least one MSNE in $\mathcal{NE}^\epsilon(\mathcal{G})$ for any $\epsilon > 0$. Using that MSNE, we can find a region surrounding it and this region is determined by the value of ϵ . We do not present a formal proof here because this is not our main focus. Instead, we concentrate on learning parameters for games with $\mu(\mathcal{NE}^\epsilon(\mathcal{G})) > 0$ for $\epsilon > 0$. Therefore, ϵ -MSNE is not only a natural and reasonable solution concept to use for the type of data captured by the CDC data but also it is convenient. The derivation of our proposed framework went through smoothly only after we solved those potential pitfalls. Indeed, we can derive technical results analogous to those of Honorio and Ortiz [16] but in our context; we present those that are most relevant to this paper. For Eqn. 1 to be valid, if $\epsilon = 1$ or $\mathcal{NE}^\epsilon(\mathcal{G}) = [0, 1]^n$, then we need to require that $q = 1$. Note that $\mathcal{NE}^\epsilon(\mathcal{G}) \neq \emptyset$ since every game has at least one MSNE [24].

3.2 Learning Games via Maximum-Likelihood

In this section, we present a way to infer games from behavioral data on mixed strategies. Let $\pi^\epsilon(\mathcal{G})$ be the *true proportion of ϵ -MSNE* in the game \mathcal{G} where $\pi^\epsilon(\mathcal{G}) \equiv \mu(\mathcal{NE}^\epsilon(\mathcal{G}))$. Given a dataset $D = \{x^{(1)}, \dots, x^{(m)}\}$, where each $x^{(l)} \sim f_{q, \mathcal{G}, \epsilon}$, i.i.d., let $\hat{\pi}^\epsilon(\mathcal{G})$ be the *empirical proportion of ϵ -MSNE*: $\hat{\pi}^\epsilon(\mathcal{G}) \equiv \frac{1}{m} \sum_{l=1}^m \mathbb{1}[x^{(l)} \in \mathcal{NE}^\epsilon(\mathcal{G})]$. Recall the *Kullback-Leibler (KL) divergence* between two Bernoulli distributions with parameters $p_1, p_2 \in (0, 1)$, which, following standard practice, we denote by $\text{KL}(p_1 \| p_2) \equiv p_1 \log \frac{p_1}{p_2} + (1 - p_1) \log \frac{1-p_1}{1-p_2}$. The following definition will be useful.

Definition 3.3. (Trivial/Non-trivial Games) A game \mathcal{G} is *trivial* iff $\mu(\mathcal{NE}^\epsilon(\mathcal{G})) \in \{0, 1\}$ and *non-trivial* iff $\mu(\mathcal{NE}^\epsilon(\mathcal{G})) \in (0, 1)$.

PROPOSITION 3.4. (Maximum-likelihood Estimation) The tuple $(\hat{\mathcal{G}}, \hat{q}, \hat{\epsilon})$ is a maximum likelihood estimator (MLE), with respect to dataset D , for the parameters of the generative model $f_{(q, \mathcal{G}, \epsilon)}$, as defined in Eqn. 1 if and only if (iff) $\hat{q} = \hat{\pi}^\epsilon(\hat{\mathcal{G}})$, and $(\hat{\mathcal{G}}, \hat{\epsilon}) \in \arg \max_{(\mathcal{G}, \epsilon)} \text{KL}(\hat{\pi}^\epsilon(\mathcal{G}) \| \pi^\epsilon(\mathcal{G}))$.

PROOF. For simplicity, we denote $\mathcal{NE}^\epsilon = \mathcal{NE}^\epsilon(\mathcal{G})$, $\pi^\epsilon \equiv \pi^\epsilon(\mathcal{G})$, and $\hat{\pi}^\epsilon \equiv \hat{\pi}^\epsilon(\mathcal{G})$. For a non-trivial \mathcal{G} , $\log f_{(q, \mathcal{G}, \epsilon)}(x^{(l)}) = \log \frac{q}{\mu(\mathcal{NE}^\epsilon)}$ for $x^{(l)} \in \mathcal{NE}^\epsilon$ and $\log f_{(q, \mathcal{G}, \epsilon)}(x^{(l)}) = \log \frac{1-q}{1-\mu(\mathcal{NE}^\epsilon)}$ for $x^{(l)} \notin \mathcal{NE}^\epsilon$. The average log-likelihood

$$\begin{aligned} \hat{\mathbb{L}}(\mathcal{G}, q, \epsilon) &= \frac{1}{m} \sum_{l=1}^m \log f_{(q, \mathcal{G}, \epsilon)}(x^{(l)}) \\ &= \hat{\pi}^\epsilon \log \frac{q}{\mu(\mathcal{NE}^\epsilon)} + (1 - \hat{\pi}^\epsilon) \log \frac{1-q}{1-\mu(\mathcal{NE}^\epsilon)} \\ &= \hat{\pi}^\epsilon \log \frac{q}{\pi^\epsilon} + (1 - \hat{\pi}^\epsilon) \log \frac{1-q}{1-\pi^\epsilon}. \end{aligned}$$

Noting that $q = \hat{q} \equiv \hat{q}(\mathcal{G}, \epsilon) \equiv \hat{\pi}^\epsilon(\mathcal{G}) = \hat{\pi}^\epsilon$ maximizes $\hat{\mathbb{L}}(\mathcal{G}, q, \epsilon)$ for any \mathcal{G} and ϵ , we obtain $\hat{\mathbb{L}}(\mathcal{G}, \hat{q}, \hat{\epsilon}) = \hat{\pi}^\epsilon \log \frac{\hat{\pi}^\epsilon}{\pi^\epsilon} + (1 - \hat{\pi}^\epsilon) \log \frac{1-\hat{\pi}^\epsilon}{1-\pi^\epsilon} = \text{KL}(\hat{\pi}^\epsilon \| \pi^\epsilon)$. \square

Let us make a few observations that follow from the MLE expression. First, if $\epsilon \geq 1$, then $\pi^\epsilon(\mathcal{G}) = 1$ for all games \mathcal{G} , which implies then $\hat{\pi}^\epsilon(\mathcal{G}) = 1$ for all games \mathcal{G} . Hence, if $\hat{\epsilon} \geq 1$ the resulting KL value is zero, so that $\hat{\mathcal{G}}$ could be *any* game. Similarly, if $\pi^\epsilon(\hat{\mathcal{G}}) = 0$

then we have $\hat{\pi}^\epsilon(\hat{\mathcal{G}}) = 0$, so that once again the resulting KL value is zero. Hence, $\hat{\mathcal{G}}$ could be *any* game. If *any* trivial game is an MLE, then *every* game, trivial or non-trivial, is also an MLE. Therefore, we can always find non-trivial games corresponding to some MLE: there are always non-trivial game MLEs.

An informal interpretation of the MLE problem is that, assuming we can keep the true proportion of ϵ -MSNE low, the learning problem simplifies to inferring a game that captures as many mixed-strategy examples in the dataset as ϵ -MSNE, but without implicitly adding more ϵ -MSNE than it needs to. Thus, formulating the learning problem using MLE brings out the fundamental tradeoff in machine learning between model complexity and “goodness-of-fit,” despite the simplicity of our model.

In fact, in most practical machine-learning (ML) applications the MLE is typically non-unique (i.e., games with different parameter values have the same ϵ -MSNE set). Common causes are *finite data* and *non-identifiability in terms of model parameters*. This non-uniqueness of the MLE is not exclusive to game hypothesis classes. Bayesian networks, for example, have the same property in general, for exactly the same reasons. Just like there are different Bayesian networks that compactly represent the same joint probability distribution, there are different games that compactly represent the same set of ϵ -MSNE. In both cases, the model parameters are non-identifiable via MLE, regardless of the amount of data available. The standard way to deal with this multiplicity of models in ML is to invoke the *Principle of Ockham’s Razor* for model selection. In general, experts in the respective field (e.g., epidemiology) would provide the necessary bias for learning. It is important to note however that, just like Bayesian networks are identifiable in terms of their joint probability distribution, games *are* identifiable in terms of their ϵ -MSNE. In both cases, inference is sound because it results from queries computed based on the object being represented, joint probability distributions or equilibria in the case of Bayesian networks or games, respectively.

3.3 Reducing MLE to Classification

One problem with the exact KL-based formulation of the MLE presented above is dealing with $\pi^\epsilon(\mathcal{G})$, even if it is well-defined (i.e., $\mu(\mathcal{NE}^\epsilon(\mathcal{G})) > 0$). Indeed, dealing with $\pi^\epsilon(\mathcal{G})$ directly would require us to compute all ϵ -MSNE of \mathcal{G} ; computing only one ϵ -MSNE is PPAD-hard in general [4, 5].

The following lemma provides bounds on the KL divergence that will be useful in our setting.

LEMMA 3.5. *Given a non-trivial game \mathcal{G} with $0 < \pi^\epsilon(\mathcal{G}) < \hat{\pi}^\epsilon(\mathcal{G})$, we can upper and lower bound the KL divergence as*

$$-\hat{\pi}^\epsilon(\mathcal{G}) \log \pi^\epsilon(\mathcal{G}) - \log 2 < \text{KL}(\hat{\pi}^\epsilon(\mathcal{G}) \| \pi^\epsilon(\mathcal{G})) < -\hat{\pi}^\epsilon(\mathcal{G}) \log \pi^\epsilon(\mathcal{G}).$$

PROOF. For simplicity, we denote $\pi^\epsilon \equiv \pi^\epsilon(\mathcal{G})$ and $\hat{\pi}^\epsilon \equiv \hat{\pi}^\epsilon(\mathcal{G})$. From the definition of KL, we have

$$\text{KL}(\hat{\pi}^\epsilon \| \pi^\epsilon) = \hat{\pi}^\epsilon \log \frac{\hat{\pi}^\epsilon}{\pi^\epsilon} + (1 - \hat{\pi}^\epsilon) \log \frac{1 - \hat{\pi}^\epsilon}{1 - \pi^\epsilon}.$$

Noting that $0 < \pi^\epsilon < \hat{\pi}^\epsilon$, we upper bound our KL expression by

$$\text{KL}(\hat{\pi}^\epsilon \| \pi^\epsilon) < \hat{\pi}^\epsilon \log \frac{\hat{\pi}^\epsilon}{\pi^\epsilon} < \hat{\pi}^\epsilon \log \frac{1}{\pi^\epsilon}$$

because $\log \frac{1-\hat{\pi}^\epsilon}{1-\pi^\epsilon} < 0$ and $\hat{\pi}^\epsilon < 1$. For the lower bound, first note that

$$\text{KL}(\hat{\pi}^\epsilon \|\pi^\epsilon) = \hat{\pi}^\epsilon \log \frac{1}{\pi^\epsilon} + (1 - \hat{\pi}^\epsilon) \log \frac{1}{1 - \pi^\epsilon} - H(\hat{\pi}^\epsilon)$$

where $H(\hat{\pi}^\epsilon) \equiv -\hat{\pi}^\epsilon \log \hat{\pi}^\epsilon - (1 - \hat{\pi}^\epsilon) \log(1 - \hat{\pi}^\epsilon)$ is the value of Shannon’s entropy function evaluated at $\hat{\pi}^\epsilon$. We obtain the lower bound because $H(\hat{\pi}^\epsilon) \leq \log 2$, by properties of the entropy function, and $\log \frac{1}{1-\pi^\epsilon} > 0$. This shows that we can bound the KL using the upper/lower bound as claimed. \square

From the last lemma, it is easy to see that when $\pi^\epsilon(\mathcal{G})$ is “low enough,” we can obtain an approximation to the MLE by simply maximizing $\hat{\pi}^\epsilon(\mathcal{G})$ only: i.e., $\arg \max_{\mathcal{G}} \text{KL}(\hat{\pi}^\epsilon(\mathcal{G}) \|\pi^\epsilon(\mathcal{G})) \approx \arg \max_{\mathcal{G}} \hat{\pi}^\epsilon(\mathcal{G})$. We implicitly enforce the constraint that $\pi^\epsilon(\mathcal{G})$ is “low enough” via regularization or some other way that allows us to introduce bias into the model selection. Therefore, we aim to develop techniques to maximize the number of ϵ -MSNE in the data while keeping ϵ small.

4 APPLICATION: GENERALIZED IDS GAMES

Given the CDC State vaccination data, we want to learn a game that would explain the behavior of the state-agents and how the behavior of the state-agents, possibly indirectly, affect the joint behavior of other state-agents within the U.S. Therefore, we want to look at reasonable games to effectively model such interaction.

Generalized IDS (α -IDS) games are well-studied games to model the investment decisions of agents when facing transfer risks (i.e., epidemics) from other agents (states). As discussed in [3, 14], α -IDS games have applicability in fire protection [19], and in vaccination settings [15]. We refer the reader to [21] for a recent survey on the broader concept of interdependent security.

In the vaccination setting, each agent decides whether or not to “invest in vaccination” (i.e., state’s effort in enforcing its population to vaccinate) given (1) the agent’s implicit and explicit cost of vaccination/effort and loss from illness of its population, (2) the vaccination decisions of other state-agents, and (3) the potential transfer probabilities/risks from other state-agents. The CDC state vaccination data captures the investment behavior each State through the vaccination rates, but does not explicitly contain the costs or losses of any state, nor the transfer risks between states. Moreover, it does not include the “average” costs, losses, or transfer risks. Below, we propose an approach to learn models involving them that best represent the CDC data at the State level.

4.1 Background

In α -IDS games with n state-agents, each state-agent i determines whether or not to invest in protection (against epidemics). Therefore, there are two actions i can play, and we denote $a_i = 1$ if i invests and $a_i = 0$ if i does not invest. We let $a = (a_1, \dots, a_n)$ to be the joint-action profile of all agents and a_{-S} to be the joint-action profile of all agents that are not in S , where for simplicity, if $S = \{i\}$, $a_{-S} \equiv a_{-i}$. There is a cost of investment C_i and loss L_i associated with the bad event occurring, either through a direct or indirect (transferred) contamination. We denote by p_i the probability that agent i will experience the bad event from a direct contamination and by q_{ji} the probability that agent i will experience the bad event due to transfer/exposure from agent j (i.e., the probability that agent j will transfer/spread the contamination/epidemics to i). Moreover,

the parameter $\alpha_i \in [0, 1]$ specifies the probability that agent i ’s investment will not protect i against transfers of a bad event. Given the parameters, the *cost function of agent i* is

$$M_i(a_i, a_{-i}) \equiv a_i[C_i + \alpha_i r_i(a_{-i})L_i] + (1 - a_i)[p_i + (1 - p_i)r_i(a_{-i})]L_i,$$

where $r_i(a_{-i}) \equiv 1 - s_i(a_{-i})$ and $s_i(a_{-i}) \equiv \prod_{j \neq i} (a_j + (1 - a_j)(1 - q_{ji}))$ are i ’s overall risk and safety functions, respectively. We aim to learn an α -IDS game from a given set of observed mixed-strategy profiles, which contain hopefully many, but not necessarily all, ϵ -MSNE. We need look at the cost function of the agents in terms of mixed-strategies. Roughly speaking, we can do this by letting x_i be the probability that $a_i = 1$ and take the expectation of the above cost function (i.e., replace all a terms by x). Comparing the cost when $a_i = 1$ and $a_i = 0$, we can derive a best-response correspondence for i . Therefore, the *expected* cost function of player i with respect to mixed-strategy x is $M_i(x_i, x_{-i}) \equiv x_i[C_i + \alpha_i r_i(x_{-i})L_i] + (1 - x_i)[p_i + (1 - p_i)r_i(x_{-i})]L_i$. By definition, an ϵ -MSNE x of an α -IDS game satisfies

$$M_i(x_i, x_{-i}) - \epsilon \leq M_i(0, x_{-i}) \text{ and } M_i(x_i, x_{-i}) - \epsilon \leq M_i(1, x_{-i}). \quad (2)$$

It follows from Eqns. 2 that $x_i[C_i + \alpha_i r_i(x_{Pa(i)})L_i] - (p_i + (1 - p_i)r_i(x_{Pa(i)}))L_i] \leq \epsilon$ and $-(1 - x_i)[C_i + \alpha_i r_i(x_{Pa(i)})L_i] - (p_i + (1 - p_i)r_i(x_{Pa(i)}))L_i] \leq \epsilon$. For simplicity, we let $\Delta_i \equiv \Delta_i(x_{Pa(i)}) \equiv C_i + \alpha_i r_i(x_{Pa(i)})L_i - (p_i + (1 - p_i)r_i(x_{Pa(i)}))L_i$.

4.2 Learning

We begin by discussing the true proportion of ϵ -MSNE to justify the use of Lemma 3.5 and then present our loss function to infer parameters for maximizing the number of ϵ -MSNE in α -IDS games.

4.3 Proportion of ϵ -MSNE in α -IDS Games

As we argue in the previous section, we can approximate our MLE objective by maximizing the number of ϵ -MSNE in the data, or equivalently, maximizing $\hat{\pi}^\epsilon(\mathcal{G})$ over ϵ and \mathcal{G} when the true proportion of ϵ -MSNE of the game is less than the empirical proportion of ϵ -MSNE in the dataset: $0 < \pi^\epsilon(\mathcal{G}) < \hat{\pi}^\epsilon(\mathcal{G})$. Below, we empirically show that the true proportion of ϵ -MSNE in α -IDS games is very low. This would justify Lemma 3.5 and our method of finding an α -IDS game that maximizes the number of ϵ -MSNE.

Fig. 1 shows the sampled proportional of randomly generated 48-player α -IDS games in various graph structures. In particular, we consider two basic graph structures that specify the transfer risks among the players. The first graph structure results from the geo-spatial adjacency of all states in the U.S.A continental (i.e., excluding Alaska and Hawaii), where each of the 48 players corresponds to a state of the US and the potential transfer risks occur from neighboring states/players. The second graph structure is based on the random graph generation of [8]. We refer to the latter type of graphs as *ER* graphs. To generate an *ER* graph, we need to specify the number of nodes and a probability $p \in [0, 1]$ that denotes the probability that the drawn *ER* graph will have an edge between any two nodes. Clearly, a higher p value corresponds to a higher density of the graph. In our case, we use *ER* graph as a way to generate different structures among the 48 players with $p \in \{0.1, 0.2, \dots, 0.9\}$. Given the graphs, we generate the values of the parameters of α -IDS games uniformly at random between zero and one. Finally, we randomly sample 100,000 mixed-strategies and check to see how many out of the 100,000 are ϵ -MSNE for $\epsilon \in \{0.1, 0.2, \dots, 0.9\}$.

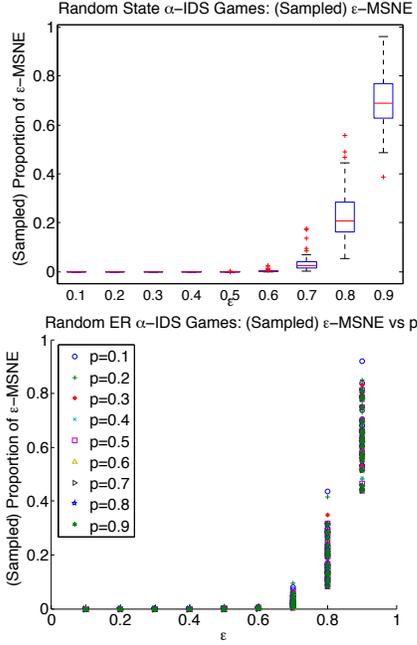


Figure 1: Sampled proportional of ϵ -MSNE in random α -IDS games. The plots show the sampled proportion of ϵ -MSNE of a fixed U.S. State topology (left) and random topology with various density (right) of random α -IDS games of 48 players. The x-axis represents the ϵ values and the y-axis represents the proportion of ϵ -MSNE.

For the random state α -IDS games, we generate 100 of them, and for each of them we compute the sampled proportion of ϵ -MSNE for each $\epsilon \in \{0.1, 0.2, \dots, 0.9\}$. Using this data, we construct the left boxplot of Fig. 1. For this plot, we observe that as ϵ goes to zero the sampled proportion of ϵ -MSNE decreases exponentially. This suggests that the true proportion of ϵ -MSNE is very small. Similarly for the random ER α -IDS games, we consider different $p \in \{0.1, 0.2, \dots, 0.9\}$, and for each fixed p , we generate 20 α -IDS games and compute the sampled proportion of ϵ -MSNE of each of them. Using this data, we construct the right boxplot of Fig. 1. We observe that as ϵ goes to zero the sampled proportion of ϵ -MSNE decreases exponentially regardless of the density and structure of the game graphs. Our results justify Lemma 3.5 and our method of finding an α -IDS game that maximizes the number of ϵ -MSNE.

4.4 Maximizing the Number of ϵ -MSNE

In our approach, we subdivide the optimization by first optimizing over \mathcal{G} , and then optimizing over ϵ . For any ϵ , we would like to apply a simple gradient-ascent optimization technique to learn the game \mathcal{G} . Unfortunately, even the latter maximization is non-trivial due to the discontinuities induced by the indicator functions $\mathbb{1}[\cdot]$ defining the ϵ -MSNE constraints. Our goal is then to further approximate $\widehat{\pi}^\epsilon(\mathcal{G})$. First, we use an upper bound by applying Eqn. 2: $\widehat{\pi}^\epsilon(\mathcal{G}) = \max_{\mathcal{G}} \frac{1}{m} \sum_{l=1}^m \mathbb{1}[x^l \in NE_\epsilon(\mathcal{G})] \leq$

$\max_{\mathcal{G}} \frac{1}{m} \sum_{l=1, i=1}^{m, n} \mathbb{1}[x_i^l \Delta_i^l \leq \epsilon] + \mathbb{1}[-(1-x_i^l) \Delta_i^l \leq \epsilon]$. Then, we approximate the indicator function in the last upper bound using a sigmoid function S , which is the standard approach leading to the BackProp algorithm in neural networks (see, e.g., [13]).

To avoid overfitting and to introduce our bias for “sparse” (graphical) game structures, we regularize the transfer parameters q_{ji} . In particular, those transfer probabilities implicitly define the structure of the α -IDS games. That is, viewing an α -IDS game from the perspective of a (directed, parametric) graphical game [18], the directed graph captures the direct transfer risks between the players where each node in the graph represents a player in the game, and there is a directed edge (i.e., an arc) from node j to node i if and only if $q_{ji} > 0$. The typical regularizer used to induce sparsity in the learned structure is the L_1 -regularizer, which we impose over the q_{ji} 's. We denote by $\lambda > 0$ the regularization parameter quantifying the amount of penalization for large values of the q_{ji} 's.

Before continuing, recall that for ϵ -MSNE to be meaningful, the cost functions should be normalized as $\widetilde{M}_i(x_i, x_{-i}) \equiv \frac{M_i(x_i, x_{-i}) - \min_i}{\max_i - \min_i}$ where $\min_i = \{C_i, p_i L_i\}$, $\max_i = \{C_i + \alpha_i r_i(0_{-i}) L_i, [p_i + (1-p_i)r_i(0_{-i})] L_i\}$, and 0_{-i} stands for the vector that sets all the elements of x_{-i} to the value 0, so that $r_i(0_{-i}) = 1 - \prod_{j \neq i} (1 - q_{ji})$. Unfortunately, working with the \widetilde{M}_i 's is cumbersome. Instead, we keep the ϵ -MSNE constraints in terms of the original unnormalized cost functions M_i 's and introduce additional constraints based on the expressions for \min_i and \max_i .

Using standard primal-dual optimization, and denoting the corresponding dual-variables/Lagrange-multipliers β_i and γ_i for each additional cost-function normalization of each player i , we obtain the following minimax program:

$$\begin{aligned} \min_{\delta, \beta} \max_{\mathcal{G}} \frac{1}{m} \sum_{l=1}^m \sum_{i=1}^n S(-x_i^l \Delta_i^l + \epsilon) + S((1-x_i^l) \Delta_i^l + \epsilon) \\ + \lambda \sum_{j=1}^n q_{ji} - \delta_i (C_i - 1) (p_i L_i - 1) - \gamma_i (2 - (C_i + \alpha_i r_i(0_{-i}) L_i)) \\ \times (2 - ([p_i + (1-p_i)r_i(0_{-i})] L_i)), \end{aligned} \quad (3)$$

where $\beta = (\beta_1, \dots, \beta_n)$ and $\gamma = (\gamma_1, \dots, \gamma_n)$. We intentionally enforce that $\min_i = \{C_i, p_i L_i\} = 1$ and $\max_i = \{C_i + \alpha_i r_i(0_{-i}) L_i, [p_i + (1-p_i)r_i(0_{-i})] L_i\} = 2$ to avoid computational issues. As long as the difference of the \min_i and \max_i is ≈ 1 , then we can easily see that the ϵ -MSNE definition will be meaningful. Our goal now is to solve the non-linear program in Eqn. 3 subject to the respective constraints on each of the variables. As stated previously, we follow the traditional approach of using gradient-ascent/descent optimization as a heuristic to update the game parameters. The process terminates when the cost functions are normalized (i.e., $\min_i = \{C_i, p_i L_i\} = 1$ and $\max_i = \{C_i + \alpha_i r_i(0_{-i}) L_i, [p_i + (1-p_i)r_i(0_{-i})] L_i\} = 2$ for every i) and after exceeding some threshold based on the total running time (i.e., ≈ 5 hours for the CDC dataset).

5 EXPERIMENTS ON VACCINATION DATA

Viewing each State as a player in the network, we interpret the vaccination percentages as mixed-strategies and generate m samples i.i.d. according to an n -variate jointly-independent Gaussian PDF, where $n = 48$, with the joint mean and standard deviations

given by each State’s reported vaccination rate and standard deviation in the CDC 2009-2010 US States H1N1 data ². For example, in the dataset, the state of Mississippi (MS) has a 17.5% vaccination rate (std. 1.5%), while that for RI is 46.8% (std. 2.0%). This is our way to account for the potential noise in the data. We consider $m \in \{500, 1000, 1500, 2000, 2500, 5000\}$. It is important to note that we do not have publicly available information about any (higher-level) correlations among states in the CDC data. Each one of the m samples is a joint mixed-strategy of dimension $n = 48$, and each component in the joint-mixed-strategy sample is drawn independently according to the mean and variance for the respective State.

5.1 Why “Obvious Baselines” Do Not Work

We are unaware of any other model/method we can apply or compare to ours. The “obvious baseline” from the perspective of probabilistic graphical models, that is, the one based on the Gaussian Markov Network representation (i.e., MRFs or Gaussian graphical models), or even Dirichlet-based graphical models, is clearly meaningless. This is because for the CDC data we only have information on the individual marginal probabilities of vaccination for each state in the U.S.: there was no covariance, or other higher-order correlation information between them available. Because our data consists of independent draws for each state from a Gaussian distribution with mean and variance as given in the data for the respective state, as the data grows, it is clear that the best probabilistic model would be a degenerate “graphical” model with no connection between the nodes: a product of Gaussians! Hence, we would have no way to infer anything about the potential interaction among the states.

Another alternatives to capture “casual non-strategic” or “strategic” relationships are structural equation models (SEMs) [25], random utility models (RUMs) [2], and quantal-response equilibrium (QRE) models [22]. In contrast to our formulation here, learning for those models boils down to a regression problem: linear (i.e., $x_i = \sum_{j \neq i} w_{ji}x_j + b_i$) for SEMs and non-linear (i.e., $x_i = S(\sum_{j \neq i} w_{ji}(2x_j - 1) + b_i)$) for RUMs and QRE. We refer to b_i as the bias term. Because we only have a single data point in the CDC dataset corresponding to a mean x_i and standard deviation σ_i for each state i , even if we were to generate data *i.i.d.* according to a $\mathcal{N}(x_i, \sigma_i^2)$, independently for each state i , as we did for learning our model, by the law of large numbers, the typical MSE regression used to learn the “obvious baselines” essentially boils down to estimating the weights using one equation (i.e. single mean) for each player. As such, the best regression model results from simply setting all weight parameters to zero (i.e., $w_{ji} = 0$ for all j) and the bias term b_i to x_i or $\ln \frac{x_i}{1-x_i}$ for SEMs or RUMs/QRE, respectively, which leads to meaningless models. If we were to remove the bias term from these models, then there would be an infinite number of weights that satisfy that single equation, and no clear way to add further constraints or preferences over the infinite solution space. For example, the obvious idea of using ridge regression with the smallest regularization parameter that would lead to unique solution yields all positive weights for the CDC dataset, which is inconsistent with the strategic substitutability one would expect in vaccination settings. Without regularization, the infinite number of solutions becomes evident

²https://www.cdc.gov/flu/fluvaxview/reportshtml/reporti0910/resources/2009-10_coverage.xlsx

as the typical MSE regression leads to very unstable values for all the aforementioned regression models, as predicted by the math. In short, these “obvious baselines” do not capture any meaningful agent interactions, nor allow us to preform strategic inferences.

As a final remark, validation is impossible in settings like ours. The same holds for other (unsupervised) settings like topic modeling, which have been found quite useful in practice, despite their only validation being the ad-hoc evaluation by means of the resulting representative most-likely words of a cluster appearing qualitatively reasonable.

5.2 Learning α -IDS Games from CDC Data

We impose an *a priori* bias for learning where only neighboring states may transfer the virus. Therefore, we are learning a geospatially-informed continental-USA State-level α -IDS game. The bias is plausible because state health departments monitor and provide neighboring states’ flu activity (i.e., www.health.state.mn.us).

To learn the parameters of an α -IDS game, we take partial derivatives of the objective function of Prog. (3) with respect to the parameters C_i, L_i, α_i, p_i , and $(q_{ji})_{j \in Pa(i)}$ for each player i and use the standard gradient-ascent optimization technique. We experiment with different regularization parameter values of β, δ, λ , and ϵ , and with various sample sizes. We present our learned α -IDS game with $\beta = \delta = -2, \lambda = 1, \epsilon = 0.35$, and $n = 1500$. We found these through empirical observations and cross-validation. The running time (to reach our termination conditions) of our algorithm increases as we increase the sample sizes. We feel that $m = 1500$ is a reasonable number within a given time limit (≈ 5 hours). From our observation, high ϵ usually results in capturing more number of ϵ -MSNE on average, while low ϵ yields the exact opposite, but *both could results in low average log-likelihood*. We use $\epsilon = 0.35$, which captures $> 90\%$ of the data as ϵ -MSNE and seems to be a good compromise.

We now present the results of the 10 learned games based on 10 different datasets generated as described above. We select a learned game among those we obtain within ≈ 5 hours with the highest accuracy for each dataset. Although the 10 games have different accuracies and log-likelihoods, they all exhibit similar behavior. We believe this indicates that our algorithm is relatively stable, and our empirical observations and conclusions are reasonably robust.

5.3 Learned α -IDS Games

Although the game parameters themselves are not our main interest, we would like to share some observations on our learned α -IDS games using the CDC H1N1 vaccination dataset because they provide anecdotal validation.

Players’ Characteristics. The first thing to note is each player’s “type.” There are two types of players in an α -IDS game, whose characterization of best-response behavior is to exhibit either *strategic complementarity* (SC) or *strategic substitutability* (SS). An SC player will play the action “vaccinate” if “enough” neighbors play action “vaccinate.” On the other hand, an SS player will play the action “vaccinate” if *not* “enough” neighbors vaccinate. Moreover, a player i is SC or SS iff $\alpha >$ or $< 1 - p$, respectively. In the vaccination setting, intuition suggest that one would expect *all* players to be SS; there is no reason for a State to vaccinate if neighboring States are protected from the virus (or epidemics). In our experiments (as

shown in Figure 2), all players in all 10 game instances we learned were SS. We did not impose any conditions on the type of players in our learning formulation: This behavior arose exclusively from properties of the CDC dataset. Such empirical results provide some partial, anecdotal evidence that the learned games are not arbitrary, and consistent with our general intuition.

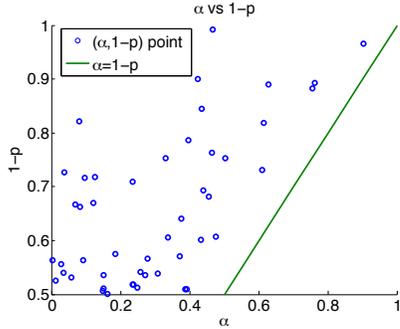


Figure 2: Players’ Types of a Learned Game. The x-axis denotes the α values of the players, the y-axis denotes the $1-p$ values of the players, and the line is the equation $\alpha = 1-p$. The plot is scaled to capture the α and $1-p$ values. The plot illustrates that our learning formulation produces values of the parameters that are consistent with vaccination scenarios, where $\alpha < 1-p$.

Player’s Best-Response Correspondences. Recall that to determine the best-response of a player, we look at the players’ best-response correspondences. In particular, the best-response correspondence of a SS player i is

$$\mathcal{BR}_i^{SS}(x_{Pa(i)}) \equiv \begin{cases} \{0\}, & \Delta_i^{SS} < s_i(x_{Pa(i)}), \\ \{1\}, & \Delta_i^{SS} > s_i(x_{Pa(i)}), \\ \{0, 1\}, & \Delta_i^{SS} = s_i(x_{Pa(i)}), \end{cases}$$

where $\Delta_i^{SS} \equiv 1 - \frac{C_i - p_i}{1 - p_i - \alpha_i}$. In order for player i to have a non-trivial response, the value of Δ_i^{SS} has to be in $(0, 1)$. Indeed, in all of our learned α -IDS games, the Δ_i^{SS} is in $(0, 1)$, for all players i . The histogram of Δ_i values of each player i fall roughly between the range of $(0.010, 0.999)$ and is appears multi-modal, with two modes near each extreme, and another near the middle of the range. Fig. 3 shows a histogram of the Δ_i values of each player i . The values fall roughly between the range of $(0.010, 0.999)$.

Players’ Transfer Risks. Recall that the transfer risks of a player are the $(q_{ji})_{j \in Pa(i)}$ where q_{ji} is the probability that a virus will transfer from j to i . Of course, our learned transfer risks depend on the mixed-strategies of the players that we use to learn the values. To show that our learned transfer risks are consistent with the training examples, we compute the safety values of each player from his neighbors using the vaccination-rate data (the mean rate we used to generate the examples). More specifically, we compute the safety value, $e_{ji} = x_j + (1-x_j)(1-q_{ji})$, of i from $j \in Pa(i)$. We also compare the values of the e_{ji} to those of values using some random mixed-strategies.

The results are shown in Fig. 4. In Fig. 4, we plot the q_{ji} and its corresponding e_{ji} values given the mean vaccination-rate (top) and

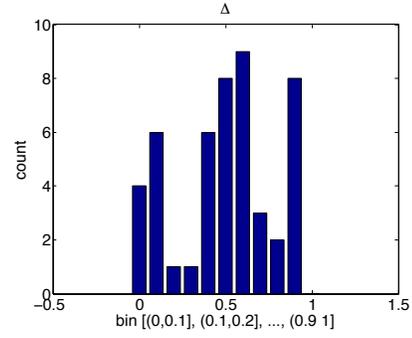


Figure 3: Histograms of the Δ values. The x-axis is discretization of the real-valued interval $[0, 1]$ into the set of intervals $\{(0, 0.1], (0.1, 0.2], \dots, (0.9, 1]\}$, while the y-axis is the count of Δ values that fall within each of those ranges.

a random mixed-strategy (bottom). The left plot shows an obvious regularity not observed on the right plot. Plotting the q_{ji} and its corresponding e_{ji} values given by the learned mean vaccination rates and those obtained from mixed-strategies drawn uniformly at random from $[0, 1]$, we observe an obvious regularity for the learned values not observed for those resulting from randomly generated mixed strategies (See Fig. 4). This suggests that the transfer risks that we learned are not random and correlated to the training examples. Hence, the results provide another piece of empirical evidence suggesting that our learned models are not arbitrary, and that, on the contrary, they seem consistent with our general intuition regarding real-world vaccination settings

Players’ Equilibrium Behavior. Our main interest for learning games is the ability they provide to potentially interpret what would happen at an MSNE, even when the given data may not consists of all examples being exact MSNE, or may be noisy. We want to infer and study the behavior of the players (i.e., US states), at an approximate MSNE of the learned game model, from noisy data, in which not all examples may belong to the set of ϵ -MSNE of some fixed but unknown game. Thus, given the learned games, we can run a version of some learning-heuristics/regret-minimization [10], in which we use the mean vaccination rates as the initial mixed-strategy profile to compute ϵ -MSNE in these games.

Fig. 5 shows the ϵ -MSNE we obtain after the best-response-gradient dynamics whenever it converges for $\epsilon \in \{0.35, 0\}$. It turns out that the mean vaccination-rates given in the CDC data is an 0.35-MSNE of the learned game. Note that this observation is non-trivial because there is no technical *a priori* reason to expect such a result: there is nothing in our learning algorithm that enforces any such condition, and the data might have as well led our learning algorithms to yield games for which such mean vaccination-rates might not have been an 0.35-MSNE of the learned game. Moreover, we are able to find an exact MSNE which is also a PSNE after trying many initial mixed-strategies that are drawn uniformly at random for the learning heuristic. *A posteriori*, the clear existence of "free-rider" states at MSNE of the learned games provides another piece of evidence consistent with the expectations of the behavior of players in vaccination-type settings. For instance, according to our learned model, at an equilibrium, NH plays the action of not

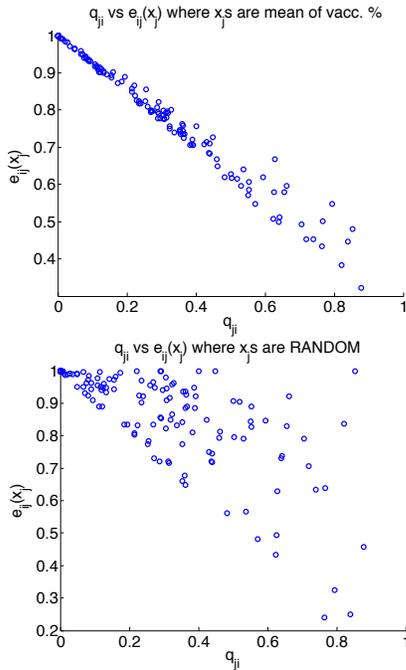


Figure 4: Values of the Safety Functions. The safety functions are evaluated using the mean of vaccination (Top) and using a random mixed-strategy (Bottom). The x-axis represents the transfer risks and the y-axis represents the values of the safety functions.

vaccinate while all of its neighbors vaccinate; we can see a similar situation for KS.

5.4 Policy Making/Analysis: An Illustration

Although the main focus of this paper is on the “learning question,” we want to show how the resulting models are useful for addressing the “inference question” more thoroughly in future work. We now provide a very brief illustration of descriptive analysis of the potential effect of increases in the vaccination rates through a targeted public-health effort in Mississippi (MS), which the CDC report as having on the lowest v_r 's for the H1N1 flu vaccine among all states in the continental U.S.A. Given the SS characteristics of our learned models one would expect that increasing the rates in one state may actually induce a lower rate on neighboring states. The question is, what is the general *magnitude* of the effect? Fig. 6 provides a possible answer to the last question within the context of our model. We see how potential interventions to increase the rate in MS, in absolute terms, affect equilibrium behavior. (Note that in any intervention, the “set” MS rate is still an ϵ -MSNE.) Increasing the rate in MS an additional 40% (almost 4-fold in relative terms) decreases that of TN by 20% in absolute terms (about 23% in relative terms). The plot suggests that increasing MS rates to about 40% may be a good compromise. The “ripple effect” on NC, a neighbor of TN but not MS, is as expected, but the magnitude is more moderate. Finally, note how the uncertainty on the possible equilibrium rates at both TN and NC increases with deliberately increasing the rate at MS.

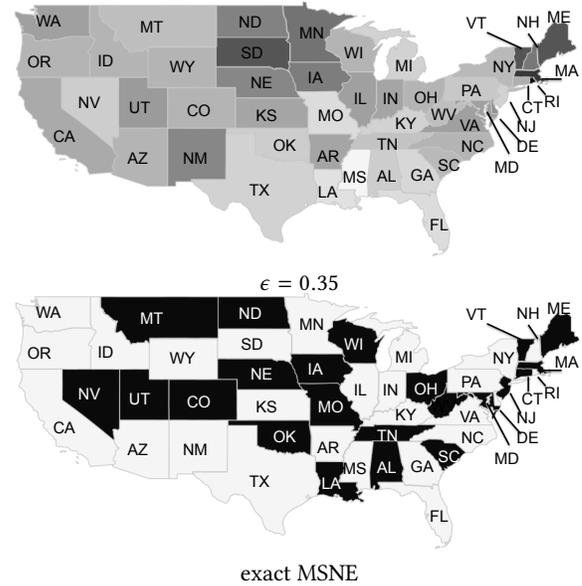


Figure 5: Equilibrium of the Learned α -IDS game. The ϵ -MSNE to which best-response-gradient dynamics consistently converged for $\epsilon \in \{0.35, 0\}$. Darker regions correspond to higher probability of vaccination (i.e., vaccination rates), for the respective ϵ -MSNE.

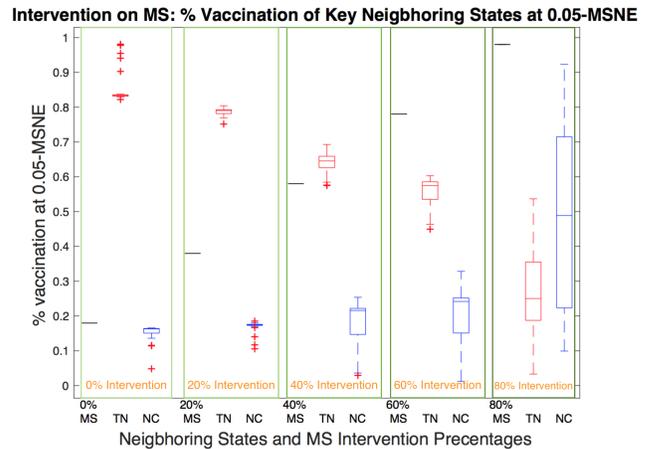


Figure 6: Studying the Potential Effect on an Intervention in MS. Approximate equilibrium vaccination rates (y-axis) in Tennessee (TN) and North Carolina (NC) as a function of potential interventions to increase the rates in Mississippi (MS) by a given percentage in absolute terms (x-axis).

6 CONCLUSION

Our interest in this work is learning games from observed mixed-strategy data. We deal with vaccination data that summarizes the actions of all the individuals within a State’s population, which can be viewed as vaccination efforts enforced by state government

officials. In our model, we view each vaccination rate as representing the mixed-strategy of each State agent, and the State agents play a variant of vaccination games among themselves. We propose and introduce a general novel ML framework to learn games from mixed-strategy data. We show how to reduce MLE to classification in our framework. We propose methods to learn α -IDS (vaccination) games from the CDC dataset. We show the effectiveness of our framework and heuristics experimentally, and illustrate policy analysis, thus providing a start to addressing this behavioral data. Our framework is general enough to learn other hypothesis class of games given a mixed-strategy dataset. While our main interest in this work is to learn vaccination games with the given CDC dataset, our framework can be applied to other applications such as learning different types of security games given the available data.

REFERENCES

- [1] Robert G. Bartle. 1995. *The Elements of Integration and Lebesgue Measure*. Wiley-Interscience.
- [2] William A. Brock and Steven N. Durlauf. 2001. Discrete Choice with Social Interactions. *The Review of Economic Studies* 68, 2 (2001), 235–260. <http://dx.doi.org/10.1111/1467-937X.00168>.
- [3] H. Chan and L. E. Ortiz. 2014. Computing Nash Equilibria in Generalized Interdependent Security Games. In *Advances in Neural Information Processing Systems 27 (NIPS'14)*. 2735–2743.
- [4] X. Chen, X. Deng, and S. Teng. 2009. Settling the Complexity of Computing Two-player Nash Equilibria. *J. ACM* 56, 3, Article 14 (May 2009), 57 pages.
- [5] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2006. The Complexity of Computing a Nash Equilibrium. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing (STOC '06)*. 71–78.
- [6] Q. Duong, Y. Vorobeychik, S. Singh, and M. P. Wellman. 2009. Learning Graphical Game Models. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*. San Francisco, CA, USA, 116–121.
- [7] Q. Duong, M. P. Wellman, S. Singh, and M. Kearns. 2012. Learning and Predicting Dynamic Networked Behavior with Graphical Multiagent Models. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (AAMAS '12)*. 441–448.
- [8] P. Erdős and A. Rényi. 1959. On random graphs, I. *Publicationes Mathematicae (Debrecen)* 6 (1959), 290–297.
- [9] S. Ficici, D. Parkes, and A. Pfeffer. 2008. Learning and Solving Many-Player Games through a Cluster-Based Representation. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI'08)*. 188–195.
- [10] D. Fudenberg and D. K. Levine. 1998. *The Theory of Learning in Games*. MIT Press Books, Vol. 1. MIT Press.
- [11] D. Fudenberg and J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.
- [12] X. A. Gao and A. Pfeffer. 2010. Learning Game Representations from Data Using Rationality Constraints. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI'10)*.
- [13] S. Haykin. 1999. *Neural Networks: A Comprehensive Foundation*. Prentice Hall.
- [14] G. Heal and H. Kunreuther. 2005. IDS Models of Airline Security. *Journal of Conflict Resolution* 49, 2 (April 2005), 201–217.
- [15] G. Heal and H. Kunreuther. 2005. *The Vaccination Game*. Working Paper. Wharton Risk Management and Decision Processes Center.
- [16] J. Honorio and L. E. Ortiz. 2015. Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *Journal of Machine Learning Research*. 16, 1 (June 2015), 1157–1210.
- [17] M. T. Irfan and L. E. Ortiz. 2014. On Influence, Stable Behavior, and the Most Influential Individuals in Networks: A Game-Theoretic Approach. *Artificial Intelligence* 215 (2014), 79–119.
- [18] M. Kearns, M. Littman, and S. Singh. 2001. Graphical Models for Game Theory. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence (UAI '01)*. 253–260.
- [19] M. Kearns and L. E. Ortiz. 2004. Algorithms for Interdependent Security Games. In *Advances in Neural Information Processing Systems (NIPS '04)*. 561–568.
- [20] M. Kearns and J. Wortman. 2008. Learning from Collective Behavior. In *In Proceedings of the 21st Annual Conference on Learning Theory (COLT '04)*. 99–110.
- [21] A. Laszka, M. Felegyhazi, and L. Buttyan. 2014. A Survey of Interdependent Information Security Games. *ACM Comput. Surv.* 47, 2, Article 23 (Aug. 2014), 38 pages.
- [22] Richard D. McKelvey and Thomas R. Palfrey. 1995. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10, 1 (1995), 6 – 38.
- [23] Jonathan S. A. Merlevede and Tom Holvoet. 2015. Game Theory and Security: Recent History and Future Directions. In *Decision and Game Theory for Security: 6th International Conference, GameSec 2015, London, UK, November 4-5, 2015, Proceedings*, MHR Khouzani, Emmanouil Panaousis, and George Theodorakopoulos (Eds.). Springer International Publishing, Cham, 334–345. https://doi.org/10.1007/978-3-319-25594-1_19
- [24] J. Nash. 1951. Non-Cooperative Games. *Annals of Mathematics* 54 (September 1951), 286–295. Issue 2.
- [25] Judea Pearl. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- [26] S. Roy, C. Ellis, S. Shiva, D. Dasgupta, V. Shandilya, and Q. Wu. 2010. A Survey of Game Theory as Applied to Network Security. In *Proceedings of the 2010 43rd Hawaii International Conference on System Sciences (HICSS '10)*. 1–10.
- [27] J. von Neumann and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press. Second Edition, 1947.
- [28] Y. Vorobeychik, M. P. Wellman, and S. Singh. 2007. Learning Payoff Functions in Infinite Games. *Mach. Learn.* 67, 1-2 (May 2007), 145–168.
- [29] J. Wright and K. Leyton-Brown. 2010. Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games. In *Proceedings of the National Conference on Artificial Intelligence (AAAI '10)*. 901–907.
- [30] J. Wright and K. Leyton-Brown. 2012. Behavioral Game Theoretic Models: A Bayesian Framework for Parameter Analysis. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (AAMAS '12)*. 921–930.
- [31] B. D. Ziebart, J. A. Bagnell, and A. K. Dey. 2010. Modeling Interaction via the Principle of Maximum Causal Entropy. In *Proceedings of the 27th International Conference on Machine Learning (ICML '10)*. 1255–1262.