PAC Algorithms for Detecting Nash Equilibrium Play in Social Networks: From Twitter to Energy Markets

WILLIAM HOILES[1], VIKRAM KRISHNAMURTHY[2], (Fellow, IEEE), AND ANUP APREM[1]
[1]Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada
[2]Department of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14850, USA
Corresponding author: A. Aprem (aaprem@ece.ubc.ca)

ABSTRACT The detection of agents whose responses satisfy equilibrium play is useful for predicting the dynamics of information propagation in social networks. Using Afriat's theorem of revealed preferences, we construct a non-parametric detection test to detect if the responses of a group of agents is consistent with play from the Nash equilibrium of a concave potential game. For agents that satisfy the detection test, it is useful to learn the associated concave potential function of the game. In this paper, a non-parametric learning algorithm is provided to estimate the concave potential function of agents with necessary and sufficient conditions on the response class for the algorithm to be a *probably approximately correct* learning algorithm. In the case of response signals measured in noise, a statistical test to detect agents playing a concave potential game that has a pre-specified Type-I error probability is provided. The detection tests and learning algorithm are applied to real-world data sets from the Twitter social network and the Ontario power grid.

INDEX TERMS Social network, Afriat's theorem, detecting equilibrium play, intertemporal utility, Twitter, energy market, PAC, revealed preferences.

I. INTRODUCTION

This paper presents a non-parametric test to detect if the response of agents is the result of equilibrium play from a concave potential game. The main question addressed is: Can the external influence and response signals from agents in a social network be used to detect if the agents are maximizing a utility function that reveals their preferences? More generally, is a dataset from a multi-agent system consistent with play from a Nash equilibrium? If yes, can the associated behavior of the agents be learned? These are important questions when considering the diffusion of information over a social network. Consider that there are two important models for information diffusion in social networks [1]: influence (contagion) based diffusion and homophily-driven diffusion. Influence-based diffusion involves agents influencing other agents to perform an action such as adopting a new technology of which several models exist to describe the dynamics [2]. Homophily-driven diffusion is based on the characteristic behavior of similar agents in the social network.[1] Detection

of agents that show common behavioral characteristics can be performed using such methods as matched sample estimation [1]. Using the tools of revealed preferences from the economics literature, this paper seeks to detect if agents in a social network share common behavioral characteristics–or equivalently, to detect if the agents have similar preferences. If agents share similar preferences, then a learning algorithm is provided to learn the associated preferences of the agents of importance for predicting the response of agents.

The preferences of interacting agents engaged in a concave potential are encoded in the concave potential function of the game. Potential games were introduced by [4] and are used extensively in the literature to study the strategic behavior of utility maximization agents. A classical example is the *congestion game* [5]–[7] in which the utility of each player depends on the amount of resource it and other players use. Recently the analysis of energy use scheduling and demand side management schemes in the energy market was performed using potential games [8]–[10]. In the energy market, the external influence is typically the price of using a particular resource, and the response is the amount of the resource used by the player. Typically, game theory is used as an analysis tool to understand the interaction of multiple

---

[1]Reference [3] gives the following illustrative example: "If your friend jumped off a bridge, would you jump too?" A possible reasons for answering "yes" is that you are friends as a result of your fondness for jumping off bridges.

agents. In contrast, this paper seeks to detect if the response signals of players are consistent with the Nash equilibrium of a concave potential game and learn the associated concave potential function. The detection of equilibrium play can be used to gain insight into the dynamics of how information propagates over a social network which is essential for social media marketing and public health campaigns.

Are there algorithms available to learn the preferences of interacting agents? In classical revealed preference theory, Afriat's theorem gives a non-parametric finite sample test to decide if an agents' response to an external influence is consistent with utility maximization [11]. Varian [12], [13], following Afriat's work, provides a non-parametric learning algorithm to estimate the utility function of single agents for predicting responses. Blundell [14] proved that the class of L-income-Lipschitz response functions is PAC-learnable using the learning algorithm provided in [15]. For multiple agents interacting (i.e. players in a game), single agent tests are not suitable. Typically the study of players in a potential game involve parametric assumptions on the utility function of the players. A notable exception is the paper presented by Deb [16] in which a non-parametric detection test for players engaged in a concave potential game is developed for intra-household consumption data based on Varian's and Afriat's work [11], [13]. Using a finite dataset, estimating the utility function of players in the game theory community has primarily followed a Bayesian parametric approach [17], [18]. For players that are engaged in a concave potential game, it is desirable to have a non-parametric learning algorithm to estimate the utility function of players for predicting future responses. This information can be used to detect and classify agents' preferences in a social network.

Consider an n-player game where each player $i \in \{1, 2, \ldots, n\}$ has a measurable set of actions $X^i \subseteq \mathbb{R}^m_+$ with generic element denoted $x^i \in \mathbb{R}^m_+$. The utility function of each player is $u^i : X \to \mathbb{R}$ where $X = \prod_{i=1}^n X^i = \{x = (x^1, x^2, \ldots, x^n) \in \mathbb{R}^{n \times m}_+\}$. The game is considered a *concave potential game* if there exists a concave potential function $V(x^1, x^2, \ldots, x^n)$ satisfying

$$u^i(x^i, x^{-i}) - u^i(x^j, x^{-i}) > 0$$
$$\text{iff } V(x^i, x^{-i}) - V(x^j, x^{-i}) > 0 \quad \forall x^i, x^j \in X^i. \quad (1)$$

Here, as in standard game theoretic notation, x^i denotes the response of player i, and x^{-i} the response of the remaining players. In words, in a potential game the incentive of all players to change their strategy is determined by a single potential function.

To formalize the problem of detecting if a dataset comprising of external influence and response signals are consistent with a concave potential game, let $t = 0, 1, \ldots$ denote discrete time, $x^i_t \in \mathbb{R}^m$ denote agent i's response with respect to the external influence $p_t \in \mathbb{R}^m$. At each discrete time t an agent is provided with an external influence p_t. Each agent i then selects a response x^i_t that maximizes their utility subject to the budget constraint $p'_t x^i_t \leq I^i_t$, where I^i_t is the budget of

agent i. Is it possible to use the time series of external influences and responses $\mathcal{D} = \{(p_t, x^1_t, x^2_t, \ldots, x^n_t)\}_{t \in \{1,2,\ldots,T\}}$ to detect if the agents are playing a concave potential game.

Why consider concave potential games? Deb, following Varian's and Afriat's work, shows that refutable restrictions exist for the dataset $\mathcal{D}$ to satisfy Nash equilibrium [11], [15], [16]. These refutable restrictions are however, satisfied by most $\mathcal{D}$ [16]. The detection of agents engaged in a concave potential game, and generating responses that satisfy Nash equilibrium, provide stronger restrictions on the dataset $\mathcal{D}$ [16], [19].

Main Results: This paper comprises 6 main results. Sec.II-D presents the preliminaries for detecting if the dataset $\mathcal{D}$ satisfies Nash rationality–that is, agents engaged in a concave potential game. For a dataset $\mathcal{D}$ that satisfies Nash rationality, Sec.II-D presents our first main result, namely a non-parametric learning algorithm to estimate the concave potential function of players in a game. The second main result of the paper, presented in Sec.II-E, deals with PAC learnability using an extension of the *fat shattering* dimension [20] to vector valued functions. As shown, if the potential function satisfies the *L-income-Lipschitz* condition then it is learnable by the learning algorithm in Sec.II-D. If a dataset $\mathcal{D}_{obs}$ fails the Nash rationality test, then it may be a result of noise. Sec.II-F provides our third main result, a statistical test to detect if a dataset $\mathcal{D}_{obs} = \{(p_t, y^1_t, y^2_t, \ldots, y^n_t) : t \in \{1, 2, \ldots, T\}\}$, $\mathcal{D}$ (5) corrupted by additive noise $y_i = x_i + w_i$ with w_i independent and identically distributed, originated from players in a concave potential game. The statistical test is shown to have a pre-specified Type-I error probability. The fourth main result, in Sec.III, is to provide a method to detect agents with similar preferences in a social network, and a test for detecting if the response of agents satisfies the maximization of an intertemporal utility function. The fifth main result is a stochastic model for the arrival times of retweets. The final main result of the paper, presented in Sec.IV, is to illustrate the algorithms on real-world datasets. The first dataset we consider is from the Twitter social network. We show that the stochastic arrival time of retweets follows a Birnbaum-Saunders distribution which provides valuable insight into the behavior of agents in online social networks [21]. It is illustrated that the agents in the Twitter social network satisfy utility maximization which allows their associated preferences to be learned. Additionally, the retweet behavior of agents satisfies intertemporal utility maximization as a result of the ageing effect associated with tweets [22], [23]. The second dataset we consider is the electrical power consumption preferences of different regions in the Ontario power system. If regions fail the detection test, then price based demand response initiatives are not suitable to control power consumption. If maximization behavior is present, then the learning algorithm presented in this paper can be used to estimate the utility function of agents. The estimated utility function can be used in the demand side management frameworks presented in [8]–[10] in which agents are engaged in a potential game.

II. REVEALED PREFERENCES: ARE AGENTS UTILITY MAXIMIZERS?

In this section we will use the principle of revealed preferences on datasets to determine how agents in social networks behave as a function of peer and external influence. The setup is depicted in the schematic diagram Fig.1.
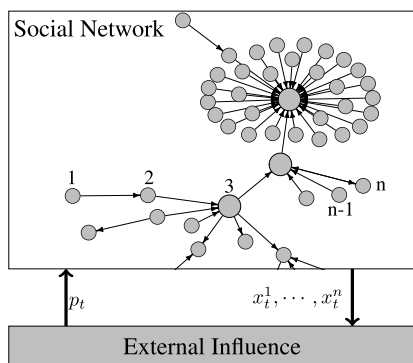

FIGURE 1. Schematic of a social network containing *n* agents where $p_t \in \mathbb{R}^m$ denotes the external influence, and $x_t^i \in \mathbb{R}^m$ the response of agent *i* in response to the external influence and other agents at time *t*. The aim is determine which agents have similar preferences.

A. AFRIAT'S THEOREM FOR NON-INTERACTING AGENTS

The theory of revealed preferences was pioneered by Samuelson [24]. The main idea behind revealed preferences is that an agent will select the response most preferred to all available responses the agent could have selected. Classical revealed preference theory has focuses on estimating the utility function of an agent based on the observed actions of the agent. Afriat's influential paper [11] in revealed preferences provides a non-parametric test to detect if the response of agents are consistent with the maximization of the agents' utility function. In this section we present Afriat's Theorem which forms the basis for detecting the preferences of interacting agents (i.e. agents engaged in a game).

Given a time-series of data $\mathcal{D} = \{(p_t, x_t), t \in \{1, 2, \ldots, T\}\}$ where $p_t \in \mathbb{R}^m$ denotes the external influence, $x_t \in \mathbb{R}^m$ denotes the response of the agent, and t denotes the time index, is it possible to detect if the agent is a *utility maximizer*? An agent is a *utility maximizer* if for every external influence p_t, the chosen response x_t satisfies

$$x_t(p_t, I_t) \in \arg\max_{\{p_t'x \le I_t\}} u(x) \qquad (2)$$

with $u(x)$ a non-satiated utility function. Non-satiated means that an increase in any element of response x results in the utility function increasing.[2] As shown by [25], without local non-satiation the maximization problem (2) may have no solution.

In (2) the budget constraint $p_t'x \le I_t$ denotes the total amount of resources available to the agent for selecting the response x to the influence p_t. For example, if $p_t = [p_t(1), \ldots, p_t(m)]$ with $p_t(i)$ the electricity price in period

$i \in m$ on day t, and $x_t = [x_t(1), \ldots, x_t(m)]$ is the associated electricity consumption in each period $i \in m$ on day t, then the budget I_t of the agent on each day t is the available monetary funds for purchasing electricity. In the real-world social network datasets provided in this paper, further insights are provided for the budget constraint.

The celebrated "Afriat's theorem" provides a necessary and sufficient condition for a finite dataset $\mathcal{D}$ to have originated from an utility maximizer.

Theorem 2.1 (Afriat's Theorem): Given a dataset $\mathcal{D} = \{(p_t, x_t) : t \in \{1, 2, \ldots, T\}\}$, the following statements are equivalent:

1) The agent is a utility maximizer and there exists a non-satiated and concave utility function that satisfies (2).
2) For scalars u_t and $\lambda_t > 0$ the following set of inequalities has a feasible solution:

$$u_\tau - u_t - \lambda_t p_t'(x_\tau - x_t) \le 0 \text{ for } t, \tau \in \{1, 2, \ldots, T\}. \qquad (3)$$

3) A non-satiated and concave utility function that satisfies (2) is given by:

$$u(x) = \min_{t \in T}\{u_t + \lambda_t p_t'(x - x_t)\} \qquad (4)$$

4) The dataset $\mathcal{D}$ satisfies the Generalized Axiom of Revealed Preference (GARP), namely for any $k \le T$, $p_t'x_t \ge p_t'x_{t+1} \quad \forall t \le k - 1 \implies p_k'x_k \le p_k'x_1$. ∎

As pointed out in Varian's influential paper [12], a remarkable feature of Afriat's theorem is that if the dataset can be rationalized by a non-trivial utility function, then it can be rationalized by a continuous, concave, monotonic utility function.

Verifying GARP (statement 4 of Theorem 2.1) on a dataset $\mathcal{D}$ comprising T points can be done using Warshall's algorithm with $O(T^3)$ [12], [26] computations. Alternatively, determining if Afriat's inequalities (3) are feasible can be done via a LP feasibility test (using for example interior point methods [27]). Note that the utility (4) is not unique and is ordinal by construction. Ordinal means that any monotone increasing transformation of the utility function will also satisfy Afriat's theorem. Geometrically the estimated utility (4) is the lower envelop of a finite number of hyperplanes that is consistent with the dataset $\mathcal{D}$.

B. AFRIAT'S THEOREM FOR INTERACTING AGENTS

We now consider a multi-agent version of Afriat's theorem for deciding if a dataset is generated by playing from the equilibrium of a potential game[3] An example is the control of power consumption in the electrical grid. Consider a corporate network of financial management operators that select the electricity prices in a set of zones in the power grid. By selecting the prices of electricity the operators are expected to be

---

[2] The non-satiated assumption rules out trivial cases such as a constant utility function which can be optimized by any response.

[3] As in [16], we consider potential games since they are sufficiently specialized so that there exist datasets that fail Afriat's test.

able to control the power consumption in each zone. The operators wish to supply their consumers with sufficient power however given the finite amount of resources the operators in the corporate network must interact. This behavior can be modelled as a game. Recent analysis of energy use scheduling and demand side management schemes in the energy market have been performed using potential games [8]–[10]. Another example of potential games are *congestion games* [4]–[7] in which the utility of each player depends on the amount of resource it and other players use.

Consider the social network presented in Fig.1, given a time-series of data from N agents $\mathcal{D} = \{(p_t, x_t^1, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}$ with $p_t \in \mathbb{R}^m$ the external influence, x_t^i the response of agent i, and t the time index, is it possible to detect if the dataset originated from agents that play a potential game?

The characterization of how agents behave as a function of external influence, for example price of using a resource, and the responses of other agents in a social network, is key for analysis. Consider the social network illustrated in Fig. 1. There are a total of n interacting agents in the network and each can produce a response x_t^i in response to the other agents and an external influence p_t. Without any *a priori* assumptions about the agents, how can the behaviour of the agents in the social network be learned? In the engineering literature the behaviour of agents is typically defined *a priori* using a *utility function*, however our focus here is on learning the behaviour of agents. The *utility function* captures the satisfaction or payoff an agent receives from a set of possible responses, denoted by X. Formally, a utility function $u : X \rightarrow \mathbb{R}$ represents a preference relation between responses x_1 and x_2 if and only if for every $x_1, x_2 \in X$, $u(x_1) \le u(x_2)$ implies x_2 is preferred to x_1. Given a time-series of data $\mathcal{D} = \{(p_t, x_t^1, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}$ with $p_t \in \mathbb{R}^m$ the influence, x_t^i the response of agent i, and t the time index, is it possible to detect if the series originated from an agent that is an *utility maximizer*?

In a social network (Fig.1), the responses of agents may be dependent on both the influence p_t and the responses of the other agents in the network, denoted by x_t^{-i}. The utility function of the agent must now include the responses of other agents–formally if there are n agents, each has a utility function $u^i(x^i, x_t^{-i})$ with x^i denoting the response of agent i, x_t^{-i} the responses of the other $n-1$ agents, and $u^i(\cdot)$ the utility of agent i. Given a dataset $\mathcal{D}$, is it possible to detect if the data is consistent with agents playing a game and maximizing their individual utilities? We denote this behaviour as *Nash rationality*, defined as follows:

Definition 2.1 ([19], [28], [29]): Given a dataset

$$\mathcal{D} = \{(p_t, x_t^1, x_t^2, \ldots, x_t^n) : t \in \{1, 2, \ldots, T\}\}, \quad (5)$$

$\mathcal{D}$ is consistent with *Nash equilibrium* play if there exist utility functions $u^i(x^i, x^{-i})$ such that

$$x_t^i(p_t, I_t^i, x^{-i}) \in \arg\max_{\{p_t' x^i \le I_t^i\}} u^i(x^i, x^{-i}). \quad (6)$$

In (6), $u^i(x, x^{-i})$ is a non-satiated utility function in x, $x^{-i} = \{x^j\}_{j \ne i}$ for $i, j \in \{1, 2, \ldots, n\}$, and the elements of p_t are strictly positive. Non-satiated means that for any $\epsilon > 0$, there exists a x^i with $\|x^i - x_t^i\|_2 < \epsilon$ such that $u^i(x^i, x^{-i}) > u^i(x_t^i, x_t^{-i})$. If for all $x^i, x^j \in X^i$, there exists a concave potential function V that satisfies

$$u^i(x^i, x^{-i}) - u^i(x^j, x^{-i}) > 0$$
$$\text{iff } V(x^i, x^{-i}) - V(x^j, x^{-i}) > 0 \quad (7)$$

for all the utility functions $u^i(\cdot)$ with $i \in \{1, 2, \ldots, n\}$, then the dataset $\mathcal{D}$ satisfies *Nash rationality*. ∎

Just as with the utility maximization budget constraint in (2), the budget constraint $p_t' x^i \le I_t^i$ in (6) models the total amount of resources available to the agents for selecting the response x_t^i to the influence p_t.

In the following sections, detection tests for utility maximization, and non-parametric learning algorithms for predicting agent responses are presented.

C. DECISION TEST FOR NASH RATIONALITY

This section presents a non-parametric test for Nash rationality given the dataset $\mathcal{D}$ defined in (5). If the dataset $\mathcal{D}$ passes the test, then it is consistent with play according to a Nash equilibrium of a concave potential game. In Sec.II-D, a learning algorithm is provided that can be used to predict the response of agents in the social network provided in Fig.1.

The following theorem provides necessary and sufficient conditions for a dataset $\mathcal{D}$ to be consistent with Nash rationality (Definition 2.1). The proof is analogous to Afriat's Theorem when the concave potential function of the game is differentiable [16], [19], [30].

Theorem 2.2 (Multiagent Afriat's Theorem): Given a dataset $\mathcal{D}$ (5), the following statements are equivalent:

1) $\mathcal{D}$ is consistent with Nash rationality (Definition 2.1) for an n-player concave potential game.

2) Given scalars v_t and $\lambda_t^i > 0$ the following set of inequalities have a feasible solution for $t, \tau \in \{1, \ldots, T\}$,

$$v_\tau - v_t - \sum_{i=1}^{n} \lambda_t^i p_t'(x_\tau^i - x_t^i) \le 0. \quad (8)$$

3) A concave potential function that satisfies (6) is given by:

$$\hat{V}(x^1, x^2, \ldots, x^n) = \min_{t \in T}\{v_t + \sum_{i=1}^{n} \lambda_t^i p_t'(x^i - x_t^i)\}. \quad (9)$$

4) The dataset $\mathcal{D}$ satisfies the Potential Generalized Axiom of Revealed Preference (PGARP) if the following two conditions are satisfied.

   a) For every dataset $\mathcal{D}_\tau^i = \{(p_t, x_t^i) : t \in \{1, 2, \ldots, \tau\}\}$ for all $i \in \{1, \ldots, n\}$ and all $\tau \in \{1, \ldots, T\}$, $\mathcal{D}_\tau^i$ satisfies GARP.

b) The responses x_t^i originated from players in a concave potential game. ∎

The intuition that connects the statements 1 and 2 in Theorem 2.2 is provided by the following results from [29]: for any smooth potential game that admits a concave potential function V, a strategy $\{x_t^1, x_t^2, \ldots, x_t^n\}$ is a pure-strategy Nash-equilibrium if and only if is a potential maximizer for each external influence p_t. Note that if only a single agent (i.e. $n = 1$) is considered, then Theorem 2.2 is identical to Afriat's Theorem. Similar to Afriat's Theorem, the constructed concave potential function (9) is ordinal–that is, unique up to positive monotone transformations.

The non-parametric test for Nash rationality involves determining if (8) has a feasible solution. Computing parameters v_t and $\lambda_t^i > 0$ in (8) involves solving a linear program with T^2 linear constraints in $(n+1)T$ variables, which has polynomial time complexity [27]. In the special case of one agent, the constraint set in (8) is the dual of the *shortest path problem* in network flows. Using the graph theoretic algorithm presented in [31], the solution of the parameters u_t and λ_t in (3) can be computed with time complexity $O(T^3)$.

D. LEARNING ALGORITHM FOR RESPONSE PREDICTION

In the previous section a non-parametric tests to detect if a dataset $\mathcal{D}$ is consistent with Nash rationality was provided. If the $\mathcal{D}$ satisfies Nash rationality, then the Multi-agent Afriat's Theorem can be used to construct the concave potential function of the game for agents in the social network illustrated in Fig.1. In this section we provide a non-parametric learning algorithm that can be used to predict the responses of agents using the constructed concave potential function (9).

To predict the response of agent i, denoted by $\hat{x}_\tau^i$, for external influence p_τ and budget I_τ^i, the optimization problem

$$\hat{x}_\tau = \{\hat{x}_\tau^1, \hat{x}_\tau^2, \ldots, \hat{x}_\tau^n\} \in \arg\max \hat{V}(\{x^i\}_{i \in \{1,2,\ldots,n\}})$$
$$\text{s.t.} \quad p_\tau' x^i \leq I_\tau^i \qquad \forall i \in \{1, 2, \ldots, n\} \qquad (10)$$

is solved using the estimated potential function $\hat{V}$ (9), p_τ, and I_τ^i. Computing $\hat{x}_\tau^i$ requires solving an optimization problem with linear constraints and concave piecewise linear objective. The optimization problem for predicting response can be formulated as a linear program and can be solved using the interior point algorithm [27]. The algorithm used to predict the response $\hat{x}_\tau = (\hat{x}_\tau^1, \hat{x}_\tau^2, \cdots, \hat{x}_\tau^n)$ is given below:

Step 1: Select an external influence $p_\tau \in \mathbb{R}_+^m$, and response budget I_τ^i for the estimation of optimal response $\hat{x}_\tau \in \mathbb{R}_+^{m \times n}$.

Step 2: For dataset $\mathcal{D}$, compute the parameters v_t and λ_t^i using (8).

Step 3: The response $\hat{x}_\tau$ is computed by solving the following linear program given $\{\mathcal{D}, p_\tau, I_\tau^i\}$, and $\{v_t, \lambda_t^i\}$

from Step 2:

$$\max \ z$$
$$\text{s.t.} \ z \leq v_t + \sum_{i=1}^{n} \lambda_t^i p_t'(\hat{x}_\tau^i - x_t^i) \text{ for } t = 1, \ldots, T$$
$$p_\tau' \hat{x}_\tau^i \leq I_\tau^i \quad \forall i \in \{1, 2, \ldots, n\} \qquad (11)$$

E. PAC-LEARNABILITY

The prediction algorithm (11) obtains a response function for the players in a concave potential game that exactly fits the dataset $\mathcal{D}$ in polynomial time. A key question is the amount of training data needed to estimate the response function–this is quantified in terms of the *sample complexity ξ*.

Definition 2.2: A class of response functions $\mathcal{C}$ from $\mathbb{R}_+^m \rightarrow \mathbb{R}_+^{m \times n}$ is *probably approximately correct* (PAC) learnable if there exists an algorithm with the following properties. For all probability distribution $\mathbb{P}$ on the external influence and response signals, $x \in \mathcal{C}$, $0 \leq \epsilon < 0.5$, and $0 \leq \delta < 1$, when the algorithm is given a dataset $\mathcal{D}$, of size $\xi = |\mathcal{D}| = \text{poly}(1/\epsilon, 1/\delta)$, outputs a hypothesis $\hat{x} \in \mathcal{C}$ in time $\text{poly}(|\mathcal{D}|)$ that satisfies the following.

 (i) $\hat{x}$ is polytime computable,
 (ii) $\hat{x}$ is consistent with x on $\mathcal{D}$,
 (iii) $\mathbb{E}_{(p,x)}(\|\hat{x}(p) - x(p)\|_\infty^2) < \epsilon$ with probability $1 - \delta$ where $x(p)$ is the actual response. ∎

In Definition 2.2, the sample complexity ξ encodes the number of training examples needed for a learner to converge with high probability to a successful hypothesis–that is, estimating the sample size necessary to ensure the hypothesis can sufficiently estimate the response for unobserved external influences. If $|\mathcal{D}| \geq \xi$, then with probability $1 - \delta$ the hypothesis $\hat{x} \in \mathcal{C}$ will be approximately correct. The learning algorithm (11) satisfies (*i*) as solving linear programs can be done in polynomial time, and from Theorem 2.2 (11) satisfies (*ii*). The following argument is used to establish that (11) satisfies (*iii*). The relationship being estimated from (11) is between external influences and responses, that is $x : \mathbb{R}_+^m \rightarrow \mathbb{R}_+^{m \times n}$. From (9), the hypothesis class $\mathcal{C}$ is the class of piecewise linear response functions. From Theorem 2.2, for any dataset satisfying Nash rationality (Definition 2.1) there is a piecewise linear response function (9) that approximates the response with zero error on $\mathcal{D}$–that is, for each external influence in $\mathcal{D}$, the piecewise linear response function can predict the associated response exactly. This implies the class $\mathcal{C}$ is dense in the class of Nash rationality responses with the supremum norm. Hence any responses that satisfy Nash rationality can be estimated with any accuracy by the class $\mathcal{C}$ of response functions.

From Definition 2.2, the sample complexity ξ must be finite for the learning algorithm (11) to be a PAC learning algorithm. Obtaining necessary and sufficient conditions for the class of response functions to be PAC-learnable by (11) therefore requires the derivation of upper and lower bounds on the sample complexity ξ. The derivation follows that

provided for utility maximization agents (responses $\mathbb{R}^m_+ \to \mathbb{R}^m_+$) in [31], and for real valued functions $\mathbb{R} \to \mathbb{R}$ in [20].

To derive the lower bound on ξ it is useful to re-define the fat shattering dimension for real valued function [20] to vector valued functions.

Definition 2.3: For $\gamma > 0$, a set of points $p_1, \ldots, p_t \in \Gamma \subset \mathbb{R}^m_+$ is γ-shattered by a class of real functions $\mathcal{C}$ if there exists $x_1, x_2, \ldots, x_T \in \mathbb{R}^{m \times n}_+$ and pairs of parallel affine hyperplanes $(H_{0,1}, H_{1,1}), \ldots, (H_{0,T}, H_{1,T})$ that satisfy the following. $H_{0,i}, H_{1,i} \subset \mathbb{R}^{m \times n}$, $\emptyset \in H^-_{0,i} \cap H^+_{1,i}$, $\text{dist}(H_{0,i}, H_{1,i}) \geq \gamma$ for $i = 1, 2, \ldots, T$. For each $b = (b_1, b_2, \ldots, b_T) \in \{0, 1\}^T$ there exists a function $\hat{x}_b \in \mathcal{C}$ such that $\hat{x}_b(p_i) \in x_i + H^+_{0,i}$ if $b_i = 0$ and $\hat{x}_b(p_i) \in x_i + H^-_{1,i}$ if $b_i = 1$–that is, the function $\hat{x}_b$ witnesses the shattering of $\{p_1, p_2, \ldots, p_t\}$. Denote $\text{fat}_{\mathcal{C}}(\gamma)$ as the fat shattering dimension defined as the maximal size of the γ-shattered set in Γ. ■

The bounds for ξ are dependent on the value of $\text{fat}_{\mathcal{C}}(\gamma)$. Hence, $\text{fat}_{\mathcal{C}}(\gamma)$ must be finite for (11) to be a PAC learning algorithm. Theorem 2.3 provides the lower bound on the sample complexity of ξ. The proof is provided in the supplementary material. The main idea of the proof is to show that there exists at least one probability distribution such that $\mathbb{E}_{(p,x)}(\|\hat{x}(p) - x(p)\|^2_\infty) > \epsilon$ if $\xi(\epsilon, \delta)$ is not sufficiently large.

Theorem 2.3: Suppose that $\mathcal{C}$ is a class of functions mapping from $\mathbb{R}^m_+ \to \mathbb{R}^{m \times n}_+$. Then any learning algorithm for $\mathcal{C}$ has sample complexity satisfying $\xi(\epsilon, \delta) \geq 0.5 \text{fat}_{\mathcal{C}}(8mn\epsilon)$.

The upper bound on the sample complexity ξ is given by Theorem 2.4. The proof for real valued functions is given in [20]; the proof for Nash rationality agents is analogous and so omitted.

Theorem 2.4: Suppose $\mathcal{C}$ is a class of response functions mapping from $\mathbb{R}^m_+ \to \mathbb{R}^{m \times n}_+$ with $\gamma < \infty$. Then any algorithm that finds function in $\mathcal{C}$ that agrees with the sample is a learning algorithm with sample complexity $\xi(\epsilon, \delta) = O\left(\frac{1}{\epsilon^2}(\ln^2(\frac{1}{\epsilon}) \text{fat}_{\mathcal{C}}(\epsilon) + \ln(\frac{1}{\delta}))\right)$ for any $\epsilon, \delta > 0$.

Theorems 2.3 and 2.4 provide the necessary and sufficient conditions for the class of response functions that can be PAC-learned by (11). As seen, the response functions must have a finite $\text{fat}_{\mathcal{C}}(\gamma)$. Unfortunately the class of all response functions has an infinite $\text{fat}_{\mathcal{C}}(\gamma)$. We consider the class of response functions that satisfy the *L-income-Lipschitz* property defined below.

Definition 2.4: ([31]) A response function $x : \mathbb{R}^m_+ \to \mathbb{R}^{m \times n}$ is *L-income-Lipschitz* if for every $p, p' \in \mathbb{R}^m_+$, there exists positive reals $L > 0$ such that

$$\frac{\|x(p) - x(p')\|_\infty}{\|p - p'\|_\infty} \leq L. \tag{12}$$

■

Note, the *L-income-Lipschitz* property does not impose any parametric assumptions on the utility function. It merely imposes a stability assumption on the responses x which

rules out dramatically different responses for similar external influences p.

The following theorem establishes that the fat shattering dimension is finite for *L-income-Lipschitz* response functions; therefore, when the dataset originates from an *L-income-Lipschitz* response function, the dataset is PAC-learnable by algorithm (11). The theorem is proved in [31] for the case $\mathbb{R}^m$ to $\mathbb{R}^m$, the extension to $\mathbb{R}^m$ to $\mathbb{R}^{m \times n}$ is analogous and therefore omitted.

Theorem 2.5: If $\mathcal{C}$ is a set of *L-income-Lipschitz* response functions from $\mathbb{R}^m_+ \to \mathbb{R}^{m \times n}$ (Definition 2.4), then the fat shattering dimension of $\mathcal{C}$ satisfies $\text{fat}_{\mathcal{C}}(\gamma) \leq (L/\gamma)^{m \times n}$ for all $L \in \mathbb{R}_+$.

F. FEASIBILITY TEST FOR NASH RATIONALITY

In real world analysis a dataset may fail the Nash rationality test (8) as a result of the response x_t being measured in noise. In this section a statistical test is provided to detect for Nash rationality in a noisy dataset, and provide a mathematical program which can be used to construct the associated concave potential function from noisy measurements.

Here we consider additive noise w_t such that measured dataset is given by:

$$\mathcal{D}_{\text{obs}} = \{(p_t, y^1_t, y^2_t, \ldots, y^n_t) : t \in \{1, 2, \ldots, T\}\}, \tag{13}$$

consisting of external influence signals p_t and noisy observations $y^i_t = x^i_t + w^i_t$. In such cases a feasibility test is required to test if the clean dataset $\mathcal{D}$ satisfies Nash rationality (8). Let H_0 and H_1 denote the null hypothesis that the clean dataset $\mathcal{D}$ satisfies Nash rationality, and the alternative hypothesis that $\mathcal{D}$ does not satisfy Nash rationality. In devising a statistical test for H_0 vs H_1, there are two possible sources of error:

> **Type-I errors:** Reject H_0 when H_0 is valid.
>
> **Type-II errors:** Accept H_0 when H_0 is invalid. (14)

Given the noisy dataset $\mathcal{D}_{\text{obs}}$ (13) the following statistical test can be used to detect if a group of agents select responses that satisfy Nash equilibrium (6) when playing a concave potential game:

$$\boxed{\int_{\Phi^*\{y\}}^{+\infty} f_M(\psi) \mathrm{d}\psi \underset{H_1}{\overset{H_0}{\gtrless}} \gamma}. \tag{15}$$

In the statistical test (15):
(i) γ is the "significance level" of the statistical test.
(ii) The "test statistic" $\Phi^*\{y\}$ is the solution of the following constrained optimization problem for $y = \{(y^1_t, y^2_t, \ldots, y^n_t)\}_{t \in \{1,2,\ldots,T\}}$:

$$\min \ \Phi$$
$$\text{s.t.} \ V_{t'} - V_t - \sum_{i=1}^{n} \lambda^i_t p_t \cdot (y^i_{t'} - y^i_t) - \sum_{i=1}^{n} \lambda^i_t \Phi \leq 0$$
$$\lambda^i_t > 0 \quad \Phi \geq 0 \quad \text{for} \ t, t' \in \{1, 2, \ldots, T\}. \tag{16}$$

(iii) f_M is the pdf of the random variable M where

$$M \equiv \underset{t,t'}{\text{Max}}[\sum_{i=1}^{n} p_t \cdot (w_t^i - w_{t'}^i)]. \qquad (17)$$

The following theorem characterizes the performance of the decision test (15). The proof is in Appendix A.

Theorem 2.6: Consider the noisy dataset $\mathcal{D}_{\text{obs}}$ (13) of external influence signals and responses. The probability that the statistical test (15) yields a Type-I error (rejects H_0 when it is true) is less then γ. (Recall H_0 and H_1 are defined in (14)). ∎

Note that (16) is non-convex due to $\sum \lambda_t^i \Phi$; however, since the objective function is given by the scalar Φ, for any fixed value of Φ, (16) becomes a set of linear inequalities allowing feasibility to be straightforwardly determined [32].

If the dataset $\mathcal{D}_{\text{obs}}$ satisfies (15), then it is desirable to estimate the associated potential function (9) which can be used to forecast response behavior of the agents using the forecasting algorithm (11) for constructing the concave potential function. Extending the results in [13] for a single agent, it is possible to estimate the lower bound on the additive errors allowing the estimation of the parameters $\{v_t, \lambda_t\}$ in (9). This involves solving the following quadratically constrained quadratic program:

$$\min \sum_{i=1}^{n} \sum_{t=1}^{T} \sum_{k=1}^{m} (y_t^i(k) - \eta_t^i(k))^2$$

$$\text{s.t. } v_t' - v_t - \sum_{i=1}^{n} \lambda_t^i p_t'[(y_\tau^i - y_t^i) - (\eta_t^i - \eta_\tau^i)] \le 0$$

$$\lambda_t^i > 0 \quad \text{for} \quad t, t' \in \{1, 2, \ldots, T\}. \qquad (18)$$

In (18) η_t denotes the minimum deviation of the observed data y_t necessary for x_t to have originated from agents satisfying Nash rationality (i.e. η_t is minimum estimate of w_t). Note that (18) will always have a feasible solution. Therefore prior to applying (18) to construct the concave potential function the dataset $\mathcal{D}_{\text{obs}}$ must pass the statistical test (15).

III. DYNAMIC MODELS OF TWITTER AGENTS

This section provides methods to cluster agents with similar preferences, a method to detect for intertemporal utility maximization, and a stochastic dynamic model for the response of agents (e.g. retweets) to tweets. We focus on the Twitter social network for the application of the methods presented in this section.

A. DETECTING AGENTS WITH SIMILAR PREFERENCES

Agents with similar preferences have a tendency to associate with alike agents, this phenomenon is known as *homophily*. Communities formed within microblogs with agents having similar race, ethnicity, age, religion, education, occupation, sex, and wealth are a direct result of homophily [1], [33], [34]. The diffusion of information over a social network is dependent on the homophily of the agents [1]. As such, being able to detect agents with similar

preferences is vital for understanding the diffusion of information in a social network. Here we present how the non-parametric utility maximization test (3) can be used to detect agents with similar preferences.

To detect for agents with similar preferences, this is identical to asking if the preference ordering of the agents are in agreement. Consider that each agent $i \in n$ has an associated dataset $\mathcal{D}^i$ of external influences and responses. Then for all the agents $i \in n$ to have similar preferences their associated preference ordering must be in agreement–that is, the collection of external influences and responses of all the agents must satisfy utility maximization (3). Formally, agents $i \in n$ with similar preferences must satisfy the following:

Definition 3.1: Given a dataset $\mathcal{D} = \{\mathcal{D}^i : i \in \{1, 2, 3, \ldots, n\}\}$ with $\mathcal{D}^i = \{(p_t^i, x_t^i) : t \in \{1, 2, \ldots, T^i\}\}$, the agents $i \in n$ have similar preferences if there exists a utility function (4) that rationalizes the dataset $\mathcal{D}$. ∎

Recall from Sec.II-C that testing if the dataset $\mathcal{D}$ satisfies utility maximization can be done using (3). Notice that the detection test for agents with similar preferences is general as the experimentalist can select the desired external influence p_t and response x_t for analysis. In real world observations it is conceivable that the response x_t is measured in noise in which case the observed dataset $\mathcal{D}_{\text{obs}}$ may fail the utility maximization test (3). In such cases the feasibility test (16) can be used to detect if the measured dataset $\mathcal{D}_{\text{obs}}$ satisfies utility maximization.

B. DETECTING INTERTEMPORAL UTILITY MAXIMIZATION

In the field of psychology and economics, the topic of intertemporal choice studies the preferences of agents over time–that is, the tradeoff between utility at different periods of time. An example is hyperbolic discounting in which agents have a preference for selecting items that arrive sooner rather then later [35], such as decreasing the time before a retweet to a tweet in the Twitter social network. Using revealed preference theory non-parametric tests for habit formation and rational anticipation are provided in [35] and [36]. In this section a non-parametric test for intertemporal choice for a measured set of external influences and responses is provided. If the dataset satisfies the test for intertemporal choice, then we illustrate how the associated intertemporal utility function can be constructed.

Consider a dataset $\mathcal{D} = \{(p_k, x_k) : k \in \{1, 2, \ldots, K\}\}$ of external influences $p_k = [p_k^1, \ldots, p_k^T] \in \mathbb{R}^{m \times T}$ and responses $x_k = [x_k^1, \ldots, x_k^T] \in \mathbb{R}^{m \times T}$ with k denoting the epoch number, T the total number of observations in each epoch, and m the number of items for each observation. A general intertemporal utility function we can consider is to test if the dataset $\mathcal{D}$ is consistent with:

$$x_k(p_k, I_k) \in \underset{\{p_k' x \le I_k\}}{\arg \max} \{U(x)\}, \qquad (19)$$

where $U(x)$ is the intertemporal utility function. A necessary and sufficient condition for the dataset $\mathcal{D}$ to satisfy (19) is provided by Afriat's Theorem. To test for

specific forms of intertemporal utility such as exponential discounting $U(x) = \sum_{t=1}^{T} \delta^{(t-1)} u(x^t)$ for $\delta \in (0, 1]$, time-separable utility $U(x) = \sum_{t=1}^{T} u^t(x^t)$, and quasi-hyperbolic discounting $U(x) = u(x^1) + \beta \sum_{t=2}^{T} u(x^t)$ for $\beta \in (0, 1]$ requires stronger restrictions than Afriat's Theorem. Here we consider the detection of *monotonic time discounting* when the utility function is given by:

$$U(x) = \sum_{t=1}^{T} D(t) u^t(x^t) \qquad (20)$$

with $D(t)$ a monotonic function $D(t) : T \to \mathbb{R}$ known as the *discounting function*, $u^t(x^t) : x^t \to \mathbb{R}$, and $U(x)$ the intertemporal utility function. The following theorem provides necessary and sufficient conditions for the dataset $\mathcal{D}$ to have originated from the intertemporal utility function (20). The proof is provided in Appendix A.

Theorem 3.1: (Intertemporal Utility): Given a dataset

$$\mathcal{D} = \{(p_k, x_k) : k \in \{1, 2, \dots, K\}\} \qquad (21)$$

the following statements are equivalent:

1) $\mathcal{D}$ is consistent with intertemporal utility maximization (19) with the non-satiated and concave utility function given by (20).

2) Given scalars u_k^t and $\lambda_k > 0$ the following set of inequalities have a feasible solution for k, $q \in \{1, \dots, K\}$ and $t \in \{1, 2, \dots, T\}$

$$u_q^t - u_k^t - \lambda_k \frac{1}{D(t)} p_k^{t\prime}(x_q^t - x_k^t) \le 0. \qquad (22)$$

3) A non-satiated intertemporal utility function that satisfies (19) is given by:

$$\hat{U}(x) = \sum_{t=1}^{T} \min_k \{D(t) u_k^t + \lambda_k p_k^{t\prime}(x^t - x_k^t)\}. \qquad (23)$$

The non-parametric test for monotonic time discounting is provided by determining if (22) has a feasible solution for a given $D(t)$. Computing the parameters u_k^t and $\lambda_k > 0$ in (22) involves solving a linear program with $K^2 T$ constraints in $(1 + T)K$ variables which has polynomial time complexity [27].

C. TWITTER RETWEET DYNAMICS

If the agents' response to a tweet satisfy intertemporal utility maximization, then it is of interest to model the arrival times of retweets. In this section a stochastic dynamic model is constructed to model the arrival times of retweets.

Consider that the retweets to a tweet arrive in a time period $[0, T]$ and the retweets are characterized by a set of moments $\{t_j^i\}(1 \le j \le n_{\text{tot}}^i)$ with n_{tot}^i the total number of retweets related to tweet i. Without loss of generality we have that $0 \le t_1^i \le \cdots \le t_{n_{\text{tot}}^i}^i \le T$. Denoting $n^i(t)$ as the total number of retweets for tweet i at time t, we consider the following model for the retweet dynamics:

$$n^i(t) = n_{\text{tot}}^i \int_0^t f^i(t; \theta^i) dt. \qquad (24)$$

In (24) $f^i(t; \theta^i)$ accounts for the temporal inhomogeneity of the retweets and is defined by a normalized probability distribution function. Note that since the retweet number is discrete valued, $n^i(t)$ is discrete valued, however for sufficiently large numbers of retweets we can take $n^i(t)$ to be continuous parameter. Typically the arrival times in social systems satisfy a Poisson process which corresponds to the arrival times satisfying an exponential distribution [37]. In the Twitter network however there may be a lag time prior to an agent reading the tweet which can be modeled using the log-normal or Birnbaum-Saunders distribution. Three possible distributions are considered in this paper: the exponential, log-normal, and Birnbaum-Saunders distributions. Using real-world data from the Twitter social network the goal is to estimate which of the three distributions is in agreement with the observed arrival times of the retweets.

Remark: To utilize (24) for predicting $n^i(t)$ requires the total number of tweets n_{tot}^i be known. In [38] a statistical relationship was constructed between the number of followers of a Twitter agent and the expected number of retweets n_{tot}^i. Denoting N_o^i as the number of followers of the agent that posted tweet i, then the following relation was empirically determined $n_{\text{tot}}^i = \alpha^i N_o^i$ with α^i generated from a log-normal distribution. The estimation of a deterministic value for n_{tot}^i is complex as the behavior of followers must be quantified. The reader is referred to [39] and [40] for further information on predicting the retweet number n_{tot}^i.

IV. REAL-WORLD DATASETS FROM TWITTER AND ENERGY MARKET

In this section we apply the decision and feasibility tests from Sec.II and Sec.III to real-world datasets from Twitter and the Ontario Energy Market. The first dataset analyzed is composed of the tweets and retweets from the Twitter Social Network. We illustrate that the retweet dynamics of tweets follows a Birnbaum-Saunders distribution, and the diffusion of information in the network is dependent on the polarity[4] of the tweets. Additionally, we illustrate that the tweeting behavior of Twitter agents satisfies utility maximization (3). The second dataset analyzed is the power consumption of different zones in the Ontario power grid. The results of the feasibility test suggest that the zones power demands are consistent with equilibrium play from a Nash equilibrium (i.e. the zones are engaged in a concave potential game over a corporate network).

---

[4]Here polarity is defined as whether the text expresses a negative-neutral-positive polarity with -1 denoting the highest negative polarity, and 1 the highest positive polarity [41].

A. TWITTER SOCIAL NETWORK: RETWEET DYNAMICS AND UTILITY MAXIMIZATION

In this section the retweet dynamics and diffusion characteristics in the Twitter network are analyzed. We show that the retweet dynamics follows a Birnbaum-Saunders distribution. With the insight gained from the analysis of the retweet dynamics and diffusion characteristics, the utility maximization test (3) is applied to investigate how the number of views and polarity affect the time before a retweet occurs and the polarity of the associated retweet. This information can be used to cluster agents with similar preferences, as discussed in Sec.III. Additionally, the intertemporal utility maximization test (22) is applied to gain insight into the Twitter agents preference for posting retweets. As discussed in [42], Twitter may rely on a huge amount of agent-generated data which can be analyzed to provide novel personal advertising to agents. Therefore results provided in this section can be used in social media marketing strategies to improve a brand and for brand awareness.

1) RETWEET DYNAMICS IN TWITTER

In this section we consider which of the three distributions (i.e. exponential, log-normal, or Birnbaum-Saunders) is in agreement with the observed retweet dynamics in the Twitter network. We also consider the importance the polarity of a tweet has on the diffusion of information in Twitter. The results provide insight for defining an appropriate external influence p_t and response x_t to test for utility maximization (3) and intertemporal utility maximization (22). The Twitter data is obtained using the *Twitter Streaming API*[5] and a custom python script. The polarity of the tweets and retweets is computed using *TextBlob*.[6] Only tweets and retweets containing the word "ebola" are considered for analysis. The complete lexical retweet network contains 4,131,521 agents with 310,709 edges. The largest connected subnetwork contains 142,236 agents and 188,903 edges and is illustrated in Fig.2 for the first 48 hours. The total duration of data collection is from October 15[th] 2014 at 9:00 pm for a duration of 290 hours.

To gain insight into how polarity and the social network impact the retweet dynamics, consider the lexical retweet network illustrated in Fig.2. The graph structure shows a tightly connected central region of highly influential agents which include Twitter accounts such as @CNN, @FoxNews, @NBCNews, and @HuffingtonPost. The edge intensity in Fig.2 provides insight on the retweet dynamics. Notice that for agents with a large *in degree* the retweets typically occur in a short period of time on the order of 1-12 hours. This behavior has been observed in popularity dynamics of literature and YouTube videos [43] and is associated with a decrease in the ability to attract new attention after ageing. This suggests that the retweet dynamics of agents may
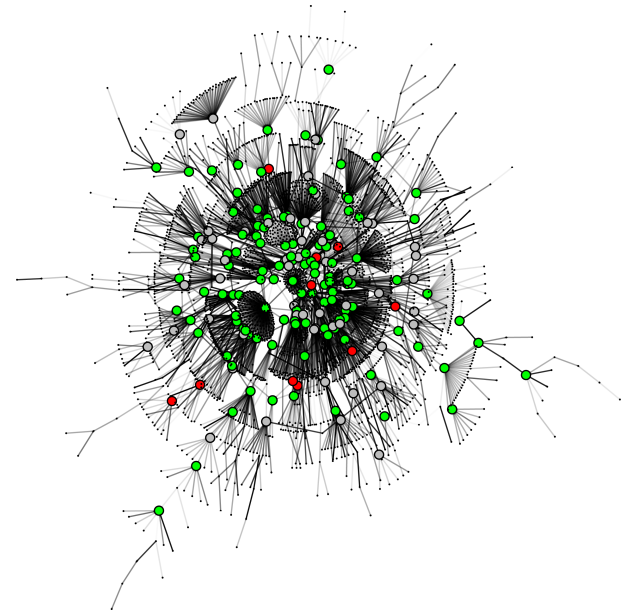
(a)
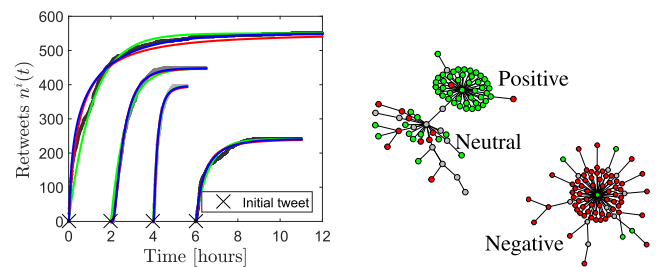
(b)

FIGURE 2. Visualization of the estimated retweet network and retweet dynamics obtained by tracking real-time tweets which contain the word "ebola" over a period of 290 hours starting from October 15[th] 2014 at 9:00 pm. Fig.2 represents the largest connected subgraph of the measured retweet network. Isolated agents (without edges) and self-loops were filtered out of the network. In Fig.2 the estimated retweet count is computed using (24) with $f^i(t; \theta^i)$ given by the exponential distribution (green line), log-normal distribution (red line), and Birnbaum-Saunders distribution (blue line). The estimated parameters θ^i are provided in Table 1. Time 0 and 6 represent tweets from @cnnbrk, time 2 from @foxnews, and at time 4 from @nytimes. (a) Lexical retweet network for "ebola". The polarity of the tweet is provided by the colour: red is negative, green positive, and grey is neutral. The polarity is only illustrated for agents with a degree larger then 5. The time of the retweets is indicated by the edge intensity, black to white with black indicating the initial time. (b) Real and estimated retweet count $n^i(t)$ for specific tweets in the Ebola retweet network in Fig.2. (c) Polarity of lexical graphs described in Sec.IV-A.

satisfy intertemporal utility maximization (22). Can the retweet dynamics of agents be modeled by (24) using an exponential, log-normal, or Birnbaum-Saunders distribution? Consider the retweet dynamics of the popular news agents @FoxNews, @cnnbrk, and @nytimes. Fig.2 provides the measured and computed $n^i(t)$ for 4 tweets (2 from @cnnbrk). As seen, the three distributions provide an accurate estimate

of the arrival times of retweets. The predictive accuracy of each is measured using the root mean squared error with the results provided in Table 1. The distribution that has the lowest root mean squared error is the Birnbaum-Saunders distribution, as seen from the results in Table 1. How does the polarity of the tweet affect the dynamics of the retweet? Let us consider two tweets for from @cnnbrk which are spaced apart by 170 hours. These are depicted in Fig.2 with the first tweet at time 0, and the second at time 6. Note that the change in followers and friends was negligible over the 170 hours between tweets, therefore only the content of the tweets impact the retweet dynamics. The tweet at time 0 has a polarity of -0.5, and at time 6 a polarity of 0.3. As seen in Fig.2, the results suggest that the greater the polarity the greater the number of retweets. Therefore the polarity of the tweets contributes to the retweet dynamics. How does the polarity of a tweet affect the polarity of retweets? Consider the lexical retweet network depicted in Fig.2. For a tweet with high polarity, both the number of retweets and the polarity of the retweets increase compared to a low polarity tweet. This suggests that increasing the polarity of a tweet will increase the polarity of the retweets.

2) UTILITY MAXIMIZATION IN THE TWITTER SOCIAL NETWORK

Using the insight gained from the analysis of the retweet dynamics, we now consider the application of the decision test (3) for utility maximization to detect for agents with similar preferences (Definition 3.1), and the detection test for intertemporal utility maximization (22) in the Twitter social network.

First we consider the detection of agents with similar preferences, as discussed in Sec.III. To apply the decision test for utility maximization (3) we must consider the definition of the external influence p_t and the associated response x_t in the Twitter network. There are two important considerations when selecting the definition of p_t and x_t. First, the agents is subject to a linear resource constraint such that $p_t' x_t = I_t$ holds. A necessary condition for this to hold is that for each $i \in m$, as $p_t(i)$ increases the associated response $x_t(i)$ must decrease. Second, the magnitude of the elements of p_t (i.e. $p_t(1), p_t(2), \ldots, p_t(m)$) should not differ by more than one order of magnitude. Additionally, the elements of x_t should not differ by more than one order of magnitude. To see the importance of the magnitude of p_t and x_t, consider the following case when $m = 2$. The associated resource budget is given by $p_t(1)x_t(1) + p_t(2)x_t(2) = I_t$. If the elements of x_t do not differ by more than one magnitude, and $p_t(1) \gg p_t(2)$, then $p_t(1)x_t(1) \approx I_t$ in which case the estimated preference of the agent is to always select $x_t(1)$. Additionally, the associated response of the agent can be estimated trivially with $x_t(1) = I_t/p_t(1)$ and $x_t(2) = 0$. To avoid this trivial case, the magnitude of the external influence and response must be considered. Here we consider the external influence to be defined by $p_t^i = [\log_{10}(\#viewers), neutrality]$ for each tweet t observed by agent $i \in n$. The *#viewers* is
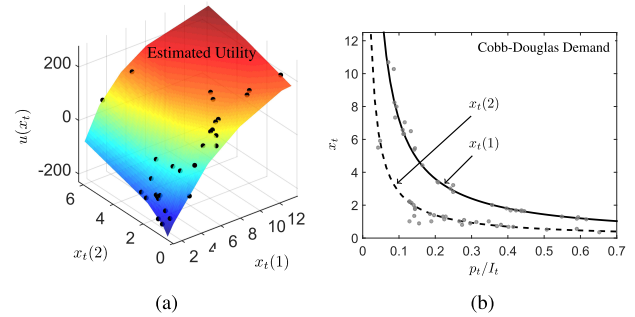

FIGURE 3. Estimated utility function $u(x_t)$, external influence and response, and the estimated Cobb-Douglas demand functions for $\mathcal{D}$ defined in Sec.IV-A. (a) Estimated utility function $u(x_t)$ using dataset $\mathcal{D}$ defined in Sec.IV-A and constructed using the non-parametric learning algorithm (4) from Afriat's Theorem. (b) Agents' response x resulting from the normalized external influence p_t/I_t—that is, $p_t' x_t/I_t = 1$. The gray dots indicate the experimental data $\mathcal{D}$ defined in Sec.IV-A, and the solid and dotted lines represent the fitted Cobb-Douglas demands $x_t(1)$ and $x_t(2)$ respectively.

estimated by the total number followers, friends, and lists the agent i contains. The *neutrality* of the tweeted text is computed as the absolute value of the reciprocal of the polarity of the tweet from TextBlob. By defining $p_t^i(1)$ as $\log_{10}(\#viewers)$ ensures that the elements of p_t^i do not differ by more than an order of magnitude. The associated response taken by each agent that retweets in the network is given by $x_t^i = [\Delta t, polarity]$ where Δt denotes the time between the tweet and retweet in minutes. The polarity of the retweet computed from TextBlob with a multiplicative factor of 10 to ensure the elements of x_t do not differ by more then one order of magnitude. Intuitively, as *#viewers* increases the time before the first retweet Δt is expected to decrease. Additionally, the results in [44] suggest that as the neutrality of the tweet decreases the associated retweet polarity is expected to increase. Therefore we consider the resource budget I_t to satisfy $I_t = p_t' x_t$.

Do all the agents in the lexical retweet network in Fig.2 have similar preferences (Definition 3.1)? Given the network is composed of $n = 142,236$ agents, the expected outcome is that all agents in the network in Fig.2 fail to have similar preferences. As expected, if we consider all agents then they fail to satisfy utility maximization and therefore do not have similar preferences. Let us consider a subgraph of agents in the lexical retweet network in Fig.2 composed of $n = 3$ agents selected at random. For these $n = 3$ agents, their responses satisfy the similar preference criterion (Definition 3.1). To gain insight into the ordinal preference relations of the $n = 3$ agents, we construct the associated utility function which is provided in Fig.3. As seen in Fig.3, agents prefer to increase the delay of retweeting compared with increasing the polarity of the retweet. Though we consider $n = 3$ agents, it is possible that other agents in Fig.2 would have similar preferences in agreement the preferences provided in Fig.3. Finding such agents involves the repeated application of (Definition 3.1) to groups of agents in Fig.2 and can be used to cluster agents with similar preferences.

Can the preferences of the $n = 3$ agents represented by the utility function in Fig.3 be estimated by the elementary

Cobb-Douglas utility function? The Cobb-Douglas utility function is given by $u(x) = \sum \alpha(i)x(i)$ with $\sum \alpha(i) = 1$. The associated Cobb-Douglas demand function is given by $x_t = \alpha(i)I_t/p)t(i)$. The parameter $\alpha(i)$ encodes the preference for the associated response $x(i)$. To test if the Cobb-Douglas utility function can be used to estimate the preferences of the agents, we fit the Cobb-Douglas demand function to the associated responses in Fig.3 by minimizing the least-square error. The results are provided in Fig.3 with $\alpha(1) = 0.72$ and $\alpha(2) = 0.28$. As expected, $x(1)$ is preferred to response $x(2)$. From the results in Fig.3, we conclude that the Cobb-Douglas utility function provides a reasonable estimate of the preferences of the $n = 3$ agents. Recall that the utility function constructed from Afrait's theorem (3) has zero error for the associated response of the agents in Fig.3.

Given the results of the retweet dynamics in Fig.2, does the retweet dynamics of agents in the Twitter social network satisfy intertemporal utility maximization (22)? The goal is to classify if the retweet dynamics satisfy a general intertemporal utility, a time-separable utility, or exponential discounting utility. Recall these intertemporal utility functions are defined below (19). The results provide new insights into the behavior of the Twitter community in response to tweets. To apply the intertemporal utility maximization test (22), we consider each tweet epoch k to be defined as the time between the first tweet and the next tweet. Each epoch k is composed of T periods. The external influence is given by $p_k^t = [p_k^t(1), p_k^t(2)]$ with $p_k^t(i)$ the total number of followers in subperiod i of period $t \in T$ in epoch k. The response is given by $x_k^t = [x_k^t(1), x_k^t(2)]$ with $x_k^t(i)$ defined as the total number of retweets in subperiod i of period $t \in T$ in epoch k. A total of $K = 20$ epochs and $T = 3$ periods are used to construct the dataset $\mathcal{D}$. Given the total number of retweets for each tweet k is $\sum_{t=1}^{T} x_k^t$, we consider the social impact budget I_k to satisfy the linear relation $I_k = p_k' x_k$. The dataset $\mathcal{D} = \{(p_k, x_k) : k \in \{1, 2, \ldots, K\}\}$ (21) is used to classify if an intertemporal utility function exists that is consistent with (20).

Using the non-parametric test (22), the dataset $\mathcal{D}$ (21) is consistent with the exponential discounted intertemporal utility function (19), with discount factor $\delta \geq 0.1$. The dataset straightforwardly satisfies the general and separable utility function since it satisfies the more restrictive detection test for exponential discounted utility. This shows that the utility obtained by a later retweet is discounted compared to that obtained for an immediate retweet. This result is expected as it has been shown for both short term and long term timescales that humans have a preference for immediate rewards compared to delayed rewards [22], [23]. As seen in Fig.2, the number of retweets dramatically increases for earlier periods of each epoch compared to later.

B. ONTARIO ELECTRICAL ENERGY MARKET
In this section we consider the application of the feasibility test for Nash rationality (15) to the aggregate power consumption of different zones in the Ontario power grid. A sampling period of $T = 79$ days starting from January 2014 is used

to generate the dataset $\mathcal{D}$ for the analysis. All price and power consumption data is available from the *Independent Electricity System Operator*[7] (IESO) website. Each zone is considered as an agent in a corporate social network. If the feasibility test is satisfied, then this suggests that zones power consumption is consistent with equilibrium play from a concave potential game. This analysis provides useful information for constructing demand side management (DSM) strategies for controlling power consumption in the electricity market. For example, if a utility function exists it can be used in the DSM strategy presented in [45] and [46]. Note that the study of corporate social networks was pioneered by Granovetter [47], [48] which shows that the social structure of the network can have important economic outcomes. Examples include agents choice of alliance partners, assumption of rational behavior, self interest behavior, and the learning of other agents behavior.

To apply the feasibility test Nash rationality (15) requires that the external influence p_t and response of agents x_t must be defined. In the Ontario power grid the wholesale price of electricity is dependent on several factors such as consumer behaviour, weather, and economic conditions. Therefore the external influence is defined as $p_t = [p_t(1), p_t(2)]$ with $p_t(1)$ the average electricity price between midnight and noon, and $p_t(2)$ as the average between noon and midnight with t denoting day. The response of each zone $i \in 10$ correspond to the total aggregate power consumption in each respective time associated with $p_t(1)$ and $p_t(2)$ and is given by $x_t^i = [x_t^i(1), x_t^i(2)]$ with $i \in \{1, 2, \ldots, n\}$. The budget I_t^i of each zone has units of dollars as p_t has units of \$/kWh and x_t^i units of kWh.

Let us first consider if the aggregate power consumption from each zone satisfies the utility maximization test (3). We find that the aggregate consumption data of each zone does not satisfy Afriat's utility maximization test (3), however, could this be a result of measurement noise? Assuming the power consumption of agents are independent and identically distributed, the central limit theorem suggests that the aggregate consumption of regions follows a zero mean normal distribution with variance σ^2. The noise term $\mathbf{w}$ in (17) is given by the normal distribution $\mathcal{N}(0, \sigma^2)$. Therefore, to test if the failure is a result of noise, the statistical test (15) is applied for each region, and the noise level σ^2 estimated for the dataset $\mathcal{D}_{obs}$ to satisfy the $\gamma = 95\%$ confidence interval for utility maximization. The results are provided in Fig.4.

As seen from Fig. 4, the Essa, West, Toronto, and East zones do not satisfy the utility maximization requirement. This results as the required noise level σ for the stochastic utility maximization test to pass is too high compared with the average power consumption. Therefore if each zone is independently maximizing then only 60% of the Ontario power grid satisfies the utility maximization test. However it is likely that the zones are engaged in a concave potential game–this would not be a surprising result as network congestion games
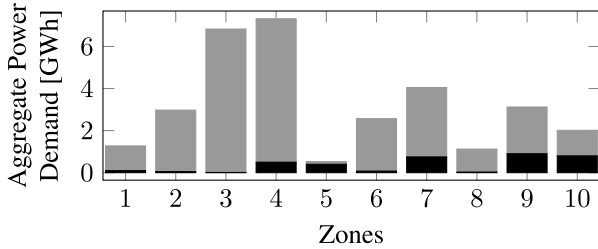
---

[7]http://ieso-public.sharepoint.com/

FIGURE 4. Average consumption (gray) and associated noise level σ (black) for the price and demand data to satisfy utility maximization in each of the 1,…,10 zones in the Ontario power grid defined in Sec.IV-B. The average hourly consumption over the $T = 79$ days starting from January 2014.

have been shown to reduce peak power demand in distributed demand management schemes [9].

To test if the dataset $\mathcal{D}$ is consistent with Nash rationality the detection test (8) is applied. The dataset for the power consumption in the Ontario power gird is consistent with Nash rationality. Using (11), a concave potential function for the game is constructed. Using the constructed potential function, when do agents prefer to consume power? The *marginal rate of substitution*[8] (MRS) can be used to determine the preferred time for power usage. Formally, the MRS of $x^i(1)$ for $x^i(2)$ is given by

$$\text{MRS}_{12} = \frac{\partial \hat{V}/\partial x^i(1)}{\partial \hat{V}/\partial x^i(2)}.$$

From the constructed potential function we find that $\text{MRS}_{12} > 1$ suggesting that the agents prefer to use power in the time period associated with $x_t(1)$–that is, the agents are willing to give up MRS_{12} kWh of power in the time period associated with $x^i(2)$ for 1 additional kWh of power in time period associated with $x^i(1)$.

The analysis in this section suggests that the power consumption behavior of agents is consistent with players engaged in a concave potential game. Using the feasibility test for Nash rationality (15) the agents preference for using power was estimated. This information can be used to improve the DSM strategies presented in [45] and [46] to control power consumption in the electricity market.

V. CONCLUSION AND FUTURE WORK
In this paper detection and feasibility tests were presented to detect if the response of agents is the result of equilibrium play from a concave potential game with a focus on learning the preferences of agents in social networks. Specifically a non-parametric feasibility test to detect if a group of agents responses are consistent with play from the Nash equilibrium of a concave potential game. To learn the preferences of agents a non-parametric learning algorithm is provided to infer the concave potential function of interacting agents engaged in a concave potential game. We prove necessary

---

[8]The amount of one good that an agent is willing to give up in exchange for another good while maintaining the same level of utility.

and sufficient conditions on the response class for the algorithm to be a *probably approximately correct* (PAC) learning algorithm. If the response signals are corrupted by noise, a statistical test to detect agents playing a game that has a pre-specified Type-I error probability is provided. The algorithms are applied to two real-world datasets from the Twitter social network and Onatrio power grid to learn the preferences of agents.

In future work we consider the refinement of the PAC learnability bounds using the fact that the demand functions are derived from monotone concave potential functions. This was not used in any derivations of the bounds for the PAC learnability requirement presented in Sec.II-E. As illustrated in [49] for single agents if the utility function is linear or linearly separable with bounded derivatives then the learning algorithm will have an improved polynomial sample-complexity compared to the general monotonic and concave utility functions considered in this paper. A standing assumption used throughout the paper is that the budget constraint is assumed to be linear. For non-linear budget sets the detection test for homophily is equivalent, however learning the associated concave potential function is non-trivial in this case. The learning algorithms in this paper rely on estimating the concave potential function from the class of piecewise-linear concave potential functions. As illustrated in [50] for single agents, utilizing (11) for prediction has the potential of generalizing very poorly. This results as the structure of external influence may bear no connection with the structure of the associated utility function. Utilizing kernel methods may offer a practical alternative to allow the underlying concave potential function to be learned in the case of a non-linear budget.

VI. APPENDIX A
PROOFS
Lemma 6.1: Suppose the functions $\{x_b : b = (b_1, b_2, \ldots, b_T) \in \{0, 1\}^T\}$ witness the shattering of $\{p_1, p_2, \ldots, p_T\}$. Then, for any $x \in \mathbb{R}_+^{m \times n}$ and labels $b, b' \in \{0, 1\}^T$ such that $b_i \neq b_i'$ for $1 \leq i \leq T$, either $\|x_b(p_i) - x\|_\infty > \gamma/(2mn)$ or $\|x_{b'}(p_i) - x\|_\infty > \gamma/(2mn)$. ∎

PROOF OF LEMMA 6.1
Proof: Since the max exceeds the mean, it follows that if x_b and $x_{b'}$ correspond to labels such that $b_i \neq b_i'$ then

$$\|x_b(p_i) - x_{b'}(p_i)\|_\infty \geq \frac{1}{mn}\|x_b(p_i) - x_{b'}(p_i)\|_2 > \frac{\gamma}{mn}.$$

This implies that for any $x \in \mathbb{R}_+^{m \times n}$ either $\|x_b(p_i) - x\|_\infty > \gamma/(2mn)$ or $\|x_{b'}(p_i) - x\|_\infty > \gamma/(2mn)$ ∎ ∎

PROOF OF THEOREM 2.3
Proof: Set $T = \frac{1}{2}\text{fat}_{\mathcal{C}}(8mn\epsilon)$, then there exists a set $P = \{p_1, p_2, \ldots, p_{2T}\}$ that is shattered by $\mathcal{C}$, the class of functions mapping from $\mathbb{R}_+^m \to \mathbb{R}_+^{m \times n}$. To prove $\xi(\epsilon, \delta) > \frac{1}{2}\text{fat}_{\mathcal{C}}(8mn\epsilon)$, it suffices to show that at least one distribution $\mathbb{P}$ of the external influences and responses in the dataset $\mathcal{D}$ (5)

requires a large sample. Equivalently, we must show that the probability of $\mathbb{E}_{(p,x)}(\|\hat{x}(p) - x(p)\|_\infty^2) > \epsilon$ is greater then zero with x a function chosen at random that witnesses the shattering of P, and $\hat{x} \in \mathcal{C}$ the hypothesis function obtained from the learning algorithm.

Denote $\mathbb{P}$ as the uniform distribution on P, and $\mathcal{C}_S = \{x_b : b = (b_1, b_2, \ldots, b_{2T}) \in \{0, 1\}^{2T}\}$ with $\mathcal{C}_S$ the set of functions that witness the shattering of P. Denote x_b as a function chosen uniformly at random from $\mathcal{C}_S$. From Lemma 6.1 with $\gamma = 8mn\epsilon$, for any $b, b' \in \{0, 1\}^{2T}$ with $b_i \neq b'_i$, either $\|\hat{x}(p) - x_b(p)\|_\infty > 4\epsilon$ or $\|\hat{x}(p) - x_{b'}(p)\|_\infty > 4\epsilon$ for any fixed function $\hat{x}$. It follows that the probability that $\|\hat{x}(p) - x_b(p)\|_\infty > 4\epsilon$ is at least 0.5. Therefore, using Markov's inequality, $\mathbb{E}_b(\|\hat{x}(p) - x_b(p)\|_\infty) > 2\epsilon$.

Consider the following randomization. Select uniformly $\iota = (i_1, i_2, \ldots, i_T) \subset [2T]$ and $b = (b_1, b_2, \ldots, b_{2T}) \in \{0, 1\}^{2T}$, and construct the random sample $S_b = \{(p_{i_j}, x_b(p_{i_j}) : j \in \{1, 2, \ldots, 2T\}\}$. Given S_b, a learning algorithm returns a function $\hat{x}_{S_b}$. It follows from Lemma 6.1, with $\gamma = 8mn\epsilon$, that the probability that $\|\hat{x}_{S_b}(p) - x_b(p)\|_\infty > 4\epsilon$ for $p \in P - \{p_{i_1}, p_{i_2}, \ldots, p_{i_T}\}$ is at least 0.5. Utilizing Markov's inequality we obtain that

$$\mathbb{E}_{b,\iota}(\mathbb{E}_{\kappa_b}(\|\hat{x}_{S_b}(p) - x_b(p)\|_\infty^2)) > 2\epsilon,$$

where κ_b is the uniform distribution on $\{(p_1, x_b(p_1)), \ldots, (p_T, x_b(p_T))\}$. Therefore, for some $b' \in \{0, 1\}^{2T}$ we have that $\mathbb{E}_\iota(\mathbb{E}_{\kappa_{b'}}(\|\hat{x}_{S_{b'}}(p) - x_{b'}(p)\|_\infty^2)) > 2\epsilon$. Without loss of generality we can assume that $\mathbb{E}_{\kappa_{b'}}(\|\hat{x}_{S_{b'}}(p) - x_{b'}(p)\|_\infty^2) > M$ for some finite M. It follows that

$$\mathbb{P}_\iota(\mathbb{E}_{\kappa_{b'}}(\|\hat{x}_{S_{b'}}(p) - x_{b'}(p)\|_\infty^2) > \epsilon) > \frac{\epsilon}{M}.$$

Hence $x_{b'}$ is not PAC-learnable by a sample of size $T = \frac{1}{2} \text{fat}_\mathcal{C}(8mn\epsilon)$. ∎

PROOF OF THEOREM 2.6

Proof: Consider a dataset $\mathcal{D}$ (5) that satisfies Nash rationality (8). Given $\mathcal{D}$, the inequalities (8) have a feasible solution. Denote the solution parameters of (8), given $\mathcal{D}$, by $\{\lambda_t^{io} > 0, V_t^o\}$. Substituting $x_t^i = y_t^i - w_t^i$, from (13), into the inequalities obtained from the solution of (8) given $\mathcal{D}$, we obtain the inequalities:

$$V_{t'}^o - V_t^o - \sum_{i=1}^n \lambda_t^{io} p_t \cdot (y_{t'}^i - y_t^i)$$
$$- \sum_{i=1}^n \lambda_t^{io} p_t \cdot (w_t^i - w_{t'}^i) \leq 0. \qquad (25)$$

Substitute $\Lambda = \sum_{i=1}^n \lambda_t^{io}$, and M, defined by (17), into (25), and recall that $\lambda_t^{io} > 0$, to obtain:

$$V_{t'}^o - V_t^o - \sum_{i=1}^n \lambda_t^{io} p_t \cdot (y_{t'}^i - y_t^i) - \Lambda M \leq 0. \qquad (26)$$

A solution of (16) given $\mathcal{D}_{\text{obs}}$, defined by (13), is denoted by $\{\Phi^*\{y\}, \lambda_t^{i*}, V_t^*\}$. By comparing the inequalities obtained

from the solution of (16) given $\mathcal{D}_{\text{obs}}$, and the inequalities (26), notice that $\{\Phi^*\{y\} = M, \lambda_t^{i*} = \lambda_t^{io}, V_t^* = V_t^o\}$ is a feasible, but not necessarily optimal solution of (16) given $\mathcal{D}_{\text{obs}}$. Therefore, for $\mathcal{D}$ satisfying malicious cooperation, it must be the case that $\Phi^*\{y\} \leq M$. This asserts, under the null hypothesis H_0, that $\Phi^*\{y\}$ is upper bounded by M. For a given $\Phi^*\{y\}$, the integral in (15) is the probability of $\Phi^*\{y\} \leq M$; therefore, the conditional probability of rejecting H_0 when true is less then γ. ∎

PROOF OF THEOREM 3.1

Proof: We prove Theorem 3.1 by construction showing that (1) → (2), (2) → (3), and (3) → (1) which provides the necessary and sufficient conditions.

(1) → (2): Suppose the dataset $\mathcal{D} = \{(p_k, x_k) : k \in \{1, 2, \ldots, K\}\}$ (21) is generated by the utility maximization (19) of a non-satiated, differentiable, and concave intertemporal utility function (20). Given the sum of a set of concave functions is concave, and that $D(t) > 0 \; \forall t \in \{1, 2, \ldots, T\}$, each utility function $u^t(x^t)$ is concave. Therefore for all $t \in \{1, 2, \ldots, T\}$ and $k, q \in \{1, 2, \ldots, K\}$ the following inequalities are satisfied:

$$u^t(x_q^t) - u^t(x_k^t) - \nabla u^t(x_k^t)'(x_q^t - x_k^t) \leq 0. \qquad (27)$$

From the Karush-Khun-Tucker conditions, if x_k solves the maximization problem (19), then there must exist Lagrange multiplies λ_k such that

$$D(t)\nabla u^t(x_k^t) = \lambda_k p_k^t \qquad (28)$$

is satisfied for all $k \in \{1, 2, \ldots, K\}$. Since $u^t(x^t)$ is a non-satiated concave function, and all the entries in p_k^t are strictly positive, the Lagrange multipliers satisfy the condition $\lambda_k > 0$. Substituting (28) into (27) the inequalities (22) result.

(2) → (3): If we have scalars u_k^t and $\lambda_k > 0$ that satisfy (22), then we must find a non-satiated and concave intertemporal utility function that rationalized the dataset $\mathcal{D}$ (21). Notice that the concavity condition provides K overestimates of the intertemporal utility function for an arbitrary bundle x since $U(x) \leq \sum_{t=1}^T D(t)[u_k^t + \lambda_k \frac{1}{D(t)} p_k^{t'}(x^t - x_k^t)]$ for $k \in \{1, 2, \ldots, K\}$. To construct the utility function we therefore take the lower envelope of the hyperplanes to obtain (23). To verify this construction works we must show that for any given x satisfying $p_k' x \leq p_k' x_k$ that $U(x) \leq U(x_k)$. First for all $k \in \{1, 2, \ldots, K\}$ we have from (23) that:

$$\max_{p_k' x \leq I_k} \{\hat{U}(x)\} \leq \max_{p_k' x \leq I_k} \{\sum_{t=1}^T D(t)u_k^t + \lambda_k p_k^{t'}(x^t - x_k^t)\}$$
$$= \sum_{t=1}^T D(t)u_k^t + \lambda_k p_k^{t'}(x_k^t - x_k^t)$$
$$= \sum_{t=1}^T D(t)u_k^t. \qquad (29)$$

TABLE 1. Fitted Distribution Parameters

	Time	0	2	4	6
Poisson	λ	1.134	0.629	0.707	0.251
	RMSE	17.81	17.70	8.698	7.268
Log-normal	σ	-0.790	-1.061	-0.611	-1.770
	μ	1.952	1.312	0.787	1.018
	RMSE	8.089	17.01	7.961	7.468
Birnbaum-Saunders	α	1.952	1.312	0.787	1.018
	β	0.367	0.334	0.539	0.164
	RMSE	6.788	16.54	6.274	6.606

Hence for all $k \in \{1, 2, \ldots, K\}$ we have that $\hat{U}(x) = \sum_{t=1}^{T} D(t) u_k^t$. For all $q \in \{1, 2, \ldots, K\}$ we have that

$$\hat{U}(x) = \sum_{t=1}^{T} \min_k \{D(t) u_k^t + \lambda_k p_k^{t\prime}(x_q^t - x_k^t)\}$$

$$= \sum_{t=1}^{T} D(t) u_q^t. \tag{30}$$

Denoting $\hat{x}_k$ as the solution of (19) given $\hat{U}(x)$ and $\mathcal{D}$, and x_k as the solution of (19) given $U(x)$ (20) and $\mathcal{D}$. Then (29) and (30) establish that $\hat{x}_k = x_k \forall k \in \{1, 2, \ldots, K\}$.

(3) $\rightarrow$ (1): Since $\hat{U}(x)$ (23) is a piecewise-linear function it also satisfies non-satiation and the concavity condition. ∎

VII. APPENDIX B
DISTRIBUTION PARAMETERS

Three distributions $f^i(t; \theta^i)$ are considered for the temporal dynamics of retweets describe by (24): the exponential, log-normal, and Birnbaum-Saunders. These functions are provided below for reference:

$$f^i(t; \lambda^i) = \lambda^i \exp(-\lambda t),$$

$$f^i(t; \sigma^i, \mu^i) = \frac{1}{\sqrt{2\pi}\,\sigma^i t} \exp\left(-\frac{(\ln(t) - \mu^i)^2}{2\sigma^{i2}}\right),$$

$$f^i(t; \alpha^i, \beta^i) = \frac{\left(\sqrt{\frac{t}{\beta^i}} + \sqrt{\frac{\beta^i}{t}}\right)}{2\alpha^i \sqrt{2\pi}\, t} \exp\left(\frac{\left(\sqrt{\frac{t}{\beta^i}} + \sqrt{\frac{\beta^i}{t}}\right)^2}{2\alpha^{i2}}\right). \tag{31}$$

Using the Twitter data provided in Fig.2, the parameters of the three distributions (31) are estimated with the results provided in Table 1. The Time in Table 1 indicates the initial time of the tweet in Fig.2, and RMSE denotes the root mean squared error.

REFERENCES

[1] S. Aral, L. Muchnik, and A. Sundararajan, "Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks," *Proc. Nat. Acad. Sci. USA*, vol. 106, no. 51, pp. 21544–21549, 2009.

[2] M. Jackson, *Social and Economic Networks*. Princeton, NJ, USA: Princeton Univ. Press, 2010.

[3] C. Shalizi and A. Thomas, "Homophily and contagion are generically confounded in observational social network studies," *Sociol. Methods Res.*, vol. 40, no. 2, pp. 211–239, 2011.

[4] D. Monderer and L. Shapley, "Potential games," *Games Econ. Behavior*, vol. 14, no. 1, pp. 124–143, 1996.

[5] R. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *Int. J. Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.

[6] A. Hayrapetyan, E. Tardos, and T. Wexler, "The effect of collusion in congestion games," in *Proc. 38th Annu. ACM Symp. Theory Comput.*, 2006, pp. 89–98.

[7] M. Babaioff, R. Kleinberg, and C. Papadimitriou, "Congestion games with malicious players," in *Proc. 8th ACM Conf. Electron. Commerce*, New York, NY, USA, 2007, pp. 103–112.

[8] A. Chapman, G. Verbič, and D. Hill, "A healthy dose of reality for game-theoretic approaches to residential demand response," in *Proc. IEEE IREP Symp. Bulk Power System Dyn. Control-IX Optim., Secur. Control Emerg. Power Grid (IREP)*, Aug. 2013, pp. 1–13.

[9] C. Ibars, M. Navarro, and L. Giupponi, "Distributed demand management in smart grid with a congestion game," in *Proc. 1st IEEE Int. Conf. Smart Grid Commun.*, Oct. 2010, pp. 495–500.

[10] C. Wu, H. Mohsenian-Rad, J. Huang, and A. Wang, "Demand side management for wind power integration in microgrid using dynamic potential game theory," in *Proc. IEEE GLOBECOM Workshops*, Dec. 2011, pp. 1199–1204.

[11] S. Afriat, "The construction of utility functions from expenditure data," *Int. Econ. Rev.*, vol. 8, no. 1, pp. 67–77, 1967.

[12] H. Varian, "The nonparametric approach to demand analysis," *Econometrica*, vol. 50, no. 1, pp. 945–973, 1982.

[13] H. Varian, "Non-parametric tests of consumer behaviour," *Rev. Econ. Stud.*, vol. 50, no. 1, pp. 99–110, 1983.

[14] R. Blundell, "How revealing is revealed preference?" *J. Eur. Econ. Assoc.*, vol. 3, nos. 2–3, pp. 211–235, 2005.

[15] H. Varian, "Revealed preference," in *Samuelsonian Economics Twenty-First Century*. Oxford, U.K.: Oxford Univ. Press, 2006, pp. 99–115.

[16] R. Deb, "Interdependent preferences, potential games and household consumption," Dept. Economics, Univ. Library Munich, Germany, MPRA Paper 6818, Jan. 2008.

[17] R. Aumann and M. Maschler, *Repeated Games with Incomplete Information*. Cambridge, MA, USA: The MIT Press, 1995.

[18] D. Fudenberg and D. Levine, *The Theory of Learning in Games*. Cambridge, MA, USA: The MIT Press, 1998.

[19] R. Deb, "A testable model of consumption with externalities," *J. Econ. Theory*, vol. 144, no. 4, pp. 1804–1816, 2009.

[20] M. Anthony and P. Bartlett, *Neural Network Learning: Theoretical Foundations*. Cambridge, U.K.: Cambridge Univ. Press, 1999.

[21] C. Leberknight, H. Inaltekin, M. Chiang, and H. Poor, "The evolution of online social networks: A tutorial survey," *IEEE Signal Process. Mag.*, vol. 29, no. 2, pp. 41–52, Mar. 2012.

[22] S. McClure, D. Laibson, G. Loewenstein, and J. Cohen, "Separate neural systems value immediate and delayed monetary rewards," *Science*, vol. 306, no. 5695, pp. 503–507, 2004.

[23] J. Kable and P. Glimcher, "The neural correlates of subjective value during intertemporal choice," *Nature Neurosci.*, vol. 10, no. 12, pp. 1625–1633, 2007.

[24] P. Samuelson, "A note on the pure theory of consumer's behaviour," *Economica*, vol. 5, no. 17, pp. 61–71, 1938.

[25] W. Diewert, "Afriat and revealed preference theory," *Rev. Econ. Stud.*, vol. 40, no. 3, pp. 419–425, 1973.

[26] A. Fostel, H. Scarf, and M. Todd, "Two new proofs of Afriat's theorem," *Econ. Theory*, vol. 24, no. 1, pp. 211–219, 2004.

[27] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.

[28] T. Ui, "Correlated equilibrium and concave games," *Int. J. Game Theory*, vol. 37, pp. 1–13, Apr. 2008.

[29] A. Neyman, "Correlated equilibrium and potential games," *Int. J. Game Theory*, vol. 26, no. 2, pp. 223–227, 1997.

[30] W. Hoiles and V. Krishnamurthy, "Nonparametric demand forecasting and detection of demand-responsive consumers," *IEEE Trans. Smart Grid*, vol. 6, no. 2, pp. 695–704, 2015.

[31] E. Beigman and R. Vohra, "Learning from revealed preference," in *Proc. 7th ACM Conf. Electron. Commerce*, New York, NY, USA, 2006, pp. 36–42.

[32] V. Krishnamurthy and W. Hoiles, "Afriat's test for detecting malicious agents," *IEEE Signal Process. Lett.*, vol. 19, no. 12, pp. 801–804, Dec. 2012.

[33] L. Kovanen, K. Kaski, and J. Kertész, and J. Saramäki, "Temporal motifs reveal homophily, gender-specific patterns, and group talk in call sequences," *Proc. Nat. Acad. Sci. USA*, vol. 110, no. 45, pp. 18070–18075, 2013.

[34] P. Block and T. Grund, "Multidimensional homophily in friendship networks," *Netw. Sci.*, vol. 2, no. 2, pp. 189–212, 2014.

[35] I. Crawford and M. Polisson, "Testing for intertemporal nonseparability," *J. Math. Econ.*, vol. 52, pp. 46–49, May 2014.

[36] I. Crawford, "Habits revealed," *Rev. Econ. Stud.*, vol. 77, no. 4, pp. 1382–1402, 2010.

[37] R. Crane and D. Sornette, "Robust dynamic classes revealed by measuring the response function of a social system," *Proc. Nat. Acad. Sci. USA*, vol. 105, no. 41, pp. 15649–15653, 2008.

[38] T. Kawamoto, "A stochastic model of tweet diffusion on the twitter network," *Phys. A, Statist. Mech. Appl.*, vol. 392, no. 16, pp. 3470–3475, 2013.

[39] Y. Bae, P. Ryu, and H. Kim, "Predicting the lifespan and retweet times of tweets based on multiple feature analysis," *ETRI J.*, vol. 36, no. 3, pp. 418–428, 2014.

[40] G. Liu, C. Shi, Q. Chen, B. Wu, and J. Qi, "A two-phase model for retweet number prediction," in *Proc. 15th Int. Conf. Web-Age Inf. Manag.*, 2014, pp. 781–792.

[41] A. Montejo-Ráez, E. Martínez-Cámara, M. Martín-Valdivia, and L. Ureña-López, "Ranked wordnet graph for sentiment polarity classification in twitter," *Comput. Speech Lang.*, vol. 28, no. 1, pp. 93–107, 2014.

[42] F. Giummolé, S. Orlando, and G. Tolomei, "A study on microblog and search engine user behaviors: How Twitter trending topics help predict Google hot queries," *Human*, vol. 2, no. 3, p. 195, 2013.

[43] H. Shen, D. Wang, C. Song, and A. Barabási. (2014). "Modeling and predicting popularity dynamics via reinforced Poisson processes." [Online]. Available: https://arxiv.org/abs/1401.0778

[44] S. Stieglitz and L. Dang-Xuan, "Emotions and information diffusion in social media-sentiment of microblogs and sharing behavior," *J. Manage. Inf. Syst.*, vol. 29, no. 4, pp. 217–248, 2013.

[45] S. Dave, M. Sooriyabandara, and L. Zhang, "Application of a game-theoretic energy management algorithm in a hybrid predictive-adaptive scenario," in *Proc. 2nd IEEE PES Int. Conf. Exhibit. Innovative Smart Grid technol. (ISGT Europe)*, Dec. 2011, pp. 1–6.

[46] H. Nguyen, J. Song, and Z. Han, "Demand side management to reduce peak-to-average ratio using game theory in smart grid," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, Mar. 2012, pp. 91–96.

[47] M. Granovetter, "Economic action and social structure: The problem of embeddedness," *Amer. J. Sociol.*, vol. 91, no. 3, pp. 481–510, 1985.

[48] G. Mark, "The impact of social structure on economic outcomes," *J. Econ. Perspect.*, vol. 19, no. 1, pp. 33–50, 2005.

[49] M. Zadimoghaddam and A. Roth, "Efficiently learning from revealed preference," in *Proc. 8th Int. Workshop, Internet Netw. Econ.*, 2012, pp. 114–127.

[50] S. Lahaie, "Kernel methods for revealed preference analysis," in *Proc. ECAI*, 2010, pp. 439–444.

WILLIAM HOILES received the M.A.Sc. degree from the Department of Engineering Science, Simon Fraser University, Burnaby, Canada, in 2012, and the Ph.D. degree from the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, Canada, in 2015. He is currently a Post-Doctoral Researcher with the Department of Electrical and Computer Engineering, The University of British Columbia. His current research interests are social sensors and the bioelectronic interface.

VIKRAM KRISHNAMURTHY (S'90–M'91–SM'99–F'05) received the Ph.D. degree from Australian National University. From 2002 to 2016, he was a Professor and the Canada Research Chair with The University of British Columbia, Canada. He is currently a Professor with the Department of Electrical and Computer Engineering, Cornell University. His research interests include statistical signal processing, computational game theory, and stochastic control in social networks. He served as a Distinguished Lecturer for the IEEE Signal Processing Society and Editor-in-Chief of the IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING. He has served as an Associate Editor for several journals including the IEEE TRANSACTIONS AUTOMATIC CONTROL and the IEEE TRANSACTIONS ON SIGNAL PROCESSING. In 2013, he was awarded an Honorary Doctorate from the Royal Institute of Technology, Sweden.

ANUP APREM received the M.E. degree from the Department of Electrical and Communication Engineering, Indian Institute of Science, Bengaluru, India, in 2012. He is currently pursuing the Ph.D. degree from the Statistical Signal Processing Laboratory, The University of British Columbia. His research interests are statistical signal processing and decision making in the area of social networks and social media.

• • •