Learnability in a Class of Non-Atomic Games arising on the Internet

Eric J. Friedman*
Department of Economics, Rutgers University
New Brunswick, NJ 08903.
friedman@econ.Rutgers.edu, *http://econ.Rutgers.edu/home/friedman*

March 6, 1998

Abstract

In this paper we study a large class of non-atomic games arising from interactions on the Internet, such as many users sharing a network link, researchers accessing a database or web server, subscribers to a network services provider trying to gain access to the modem pool and many more.

We provide simple and easily computable conditions under which these games are learnable by several models of learning, such as adaptive and sophisticated learning (Milgrom and Roberts [34]), calibrated learning (Foster and Vohra [8]), and reasonable learning (Friedman and Shenker [15]). Thus, using these methods one can evaluate the stability of such games on the Internet or in other settings in which players must learn over time.

Running head: Learnability in Internet Games

*I would like to thank Michelle Goman, Adam Landsberg, Hervé Moulin, and Scott Shenker for helpful conversations. An earlier version of this paper was circulated under the title "Non-Atomic Games with Multiple Externalities"

1 Introduction

The Internet is currently undergoing a dramatic transformation from a (relatively) small network populated by computer literate and generally cooperative users to a gigantic network with very different characteristics where behavior is much more variable and "noncooperative." Such noncooperative behavior is clearly seen in the growth of Spam-email or more subtly in the competition among browsers which can open multiple TCP connections to increase performance. Thus, it is of increasing importance to understand incentive issues on the Internet and there has been a growing body of work on this topic, e.g., [31, 37, 5, 41, 42, 24, 25, 39, 38, 4, 3, 30, 29, 36, 32].

To address these issues, many theorists have begun to model the Internet as a "noncooperative game" (see, e.g. [7, 6, 17, 19, 26, 25]). Most of these analyses assume that the outcome of a game is a Nash equilibrium, the standard outcome analyzed in economics [16]. Given the decentralized and complex structure of the Internet it is not realistic that players would just "know" (from detailed mathematical analysis) what the current Nash equilibrium was. (Note that every time a new user logs on, or a system crashes the Nash equilibrium changes.) Thus, it seems clear that if a Nash equilibrium arose it would be due to some learning process. However, this reasoning is surprisingly difficult to justify; until recently there was no known decentralized learning procedure which was guaranteed to converge to a Nash equilibrium of an arbitrary game; however the convergence rate of this process [9] is far too slow to be of practical use. Thus, the question of what the outcome of a noncooperative game will be is at present unsolved.

For many questions of interest this can be restated as an attempt to characterize the set of asymptotic plays – those strategies which are played 'often' after an initial 'transient', or learning period. Although, no reasonable model of learning has been shown which always converges to the set of Nash equilibria, many dynamic models have been shown to converge the "serially undominated set" under restrictive informational assumptions, for synchronous play (see, e.g. [34]) and more generally to the "serially unoverwhelmed set[1]" in noisy and asynchronous environments [15].

Thus, certain games are inherently learnable, while others are not. In this paper we consider a large class games which arise on the Internet for which it is easy to determine their "degree of learnability." We restrict ourselves to the class of nonatomic games, which are model games with many players.

As an example, consider a typical game arising from the Internet in which players share a resource with negative congestion effects. Instances of such games include users sending information over a network, sharing a database, accessing a web server, or dialing in to their service provider's modem pool.[2] In these games, a player who increases her use of the resource imposes a negative externality on the other players. The key aspect of many of these games is that players know of the externality (e.g. congestion and delays) but cannot quantify its affect. This is due both to the complexity of the game, since the payoff functions in these games are often unknown analytically and also are stochastically complex, and their

[1]The serially unovewhelmed set will be formally defined in Section 7. Intuitively, it arises because a learner with limited information is unable to identify dominated strategies.

[2]In economics, similar games also arise between polluting factories, fishermen using common waters, and markets with close substitutes.

inability to observe the other players strategies. Thus they must learn (or optimize) over time.

Our main results show that in many cases the analysis of these complicated set valued properties, which are the conditions for "learnability," can be reduced to the much simpler analysis of (a point valued version of) the best reply mapping.[3] This can be further reduced to the analysis of a single trajectory, which can be easily analyzed or simulated on a computer. In particular, using this result we show that many games with sufficient "capacity" (to be defined in Section 7) are learnable in both senses.

This paper is organized as follows. In the next section we discuss our learnability assumptions in more detail. Then in Section 3 we formally define an non atomic externality game (NAEG) and provide several examples in Section 4. Section 5 analyzes the properties of best reply dynamics for NAEGs and also, as a separate result of technical interest, proves a basic theorem about the structure of the equilibria for these games. Section 6 analyzes the learnability of NAEGs under iteration elimination of dominated strategies, while Section 7 extends these results to learnability under the iterated elimination of overwhelmed strategies. Section 8 concludes with several applications of these results both analytic (convergence with sufficient capacity) and numerical.

[3]Note that we do not consider best reply (Cournot) dynamics to be a reasonable model of learning; it is merely a useful tool to check for convergence of more complex learning processes.

2 Learning and Solution Concepts

In this section we informally review the literature on learning to provide motivation for our formal analysis of the two solution concepts discussed in this paper.

The first such concept, the iterated elimination of strictly dominated strategies, is well known, and was shown by Milgrom and Roberts [34] to contain the set of asymptotic plays under both adaptive and sophisticated learning. As discussed in that paper, many of the learning procedures studied in the literature are contained in this class including, best-reply (Cournot) dynamics, fictitious play, and many versions of Bayesian learning. When the iterated elimination of dominated strategies leads to a unique strategy profile, then asymptotic play under any of these procedures converges to the Nash equilibrium of the game which is guaranteed to be unique. We denote such games as D-solvable. D-solvable games also yield unique predictions under evolutionary learning [43, 20] and in synchronous[4] games with probabilistic learners [14].

From a different perspective, Foster and Vohra [8] study the behavior of calibrated learners in synchronous games. They show that a group of calibrated learners will converge to a correlated equilibrium of the underlying game. Since the set of correlated equilibria is contained in the serially undominated set, which is the outcome of the iterated elimination of dominated strategies, our results on the D-solvability of games also imply the convergence

[4] A synchronous game is one in which all players choose their strategies simultaneously in discrete time after having received payoff information about the previous period and is the standard model of a repeated game. Asynchronous games, in which players choose their strategies at different times based on time lagged information, have been less studied. In particular [27] studies the subgame perfect Nash equilibria of a class of asynchronous games and [14, 15] analyze the convergence of learning procedures in this setting.

of groups of calibrated learners.

Lastly, Friedman and Shenker [14, 15] consider the asymptotic behavior of "reasonable learning procedures" on an Internet like setting, with extremely limited information and asynchronous play. They show that such procedures are not guaranteed to converge even in D-solvable games. However, they define a new concept, the overwhelmed strategy and show that if a game is solvable by the repeated elimination of overwhelmed strategies (O-solvable) then reasonable learners will converge to the Nash equilibrium. Thus, when the payoff function is complex and not known by the players, strategies are not observable, and play is asynchronous, as on the Internet, we believe that O-solvability is the crucial property of a game.[5]

3 Non Atomic Externality Games

We define a non atomic[6] externality game (NAEG) to be a non-atomic game with m classes of players and 'submodular' payoff functions which depend only on a player's strategy and the 'externality' (to be defined). In this section we present the technical conditions which define an NAEG, which will then be demonstrated by several examples in the following section.

Let $X = X_1 \bigcup X_2 \ldots \bigcup X_m$, where each X_i is a measure space with an associated measure

[5]Recent attempts to verify this theory include Chen (1997), who did an experimental study of asynchronous games, and Friedman, Greenwald, and Shenker (1998) who have computationally simulated learning processes. The results of both papers appear to be in accord with O-solvability.

[6]A non-atomic game is a game with a continuum of players, in which no single player can affect the other players by changing only her own action. The classic reference for such games is [1]. Applications of non-atomic games to network and queuing problems include [32, 44, 10, 12, 13]. In particular, [44] provides a good introduction to this class of games.

μ_i, with $\mu_i(X_i) = N_i$ the population of class i.[7] Let the strategy set Σ be a measurable mapping from X to compact subsets of $\Re$. (Different players may have different strategy spaces, i.e. $\Sigma(x)$ is the strategy space for player x.) We will view Σ as a fiber bundle over X with fibers $\Sigma(x)$, as this will prove to be notationally convenient.[8]

A strategy profile is a measurable function $s : X \to A$, where $s(x) \in \Sigma(x)$. Note that we will denote a strategy profile by lower case letters, and often it will be convenient to think of them (with a slight abuse of notation) as subbundles of Σ with the property that $|s(x)| = 1$, for all $x \in X$. For any A which is a compact subbundle of Σ, let $\hat{\Xi}(A)$ be the set of all compact subbundles of A and let $\Xi(A)$ be the subset of $\hat{\Xi}(A)$ consisting of strategies.[9] To simplify notation we will often write Ξ for the set of all strategies as an abbreviation for $\Xi(\Sigma)$.

Define the externality of a strategy profile, $s \in \Xi$ by $\lambda(s) = (\lambda_1(s), \ldots, \lambda_m(s))$ with $\lambda_i(s) = \int_{X_i} s(x_i) d\mu_i(x_i)$, where $\lambda_i(s)$ is the total 'demand' of players in class i. It will be convenient to consider the space of externalities as a lattice, $\lambda(s) \in \Lambda = [0,1]^m$ with the standard vector ordering.

The key restriction on player x's payoff is that it only depends on her own strategy $s(x)$ and the externality vector $\lambda(s)$. Thus we define (with a slight abuse of notation)

$$\Gamma(s, x) = \Gamma_x(s(x), \lambda(s))$$

[7] Note that by definition, $x \in X$ represents a particular player from a particular class; occasionally we will emphasize this for a player in class i with the notation x_i; however, this does *not* imply that x is a vector, as x will be used to denote a (single) generic player.

[8] A fiber bundle is a space E with a base space $B \subseteq E$ and a projection $p : E \to B$ where the set $p^{-1}(b)$ is called the fiber over $b \in B$. Note that in our notation $\Sigma(x)$ is used to denote $p^{-1}(x)$. See, e.g., [18] for more details about fiber bundles.

[9] We say that a subbundle S is compact if for $\forall x \in X$ the set $S(x)$ is compact.

where Γ_x is a (measurable) function with decreasing differences: for all $\sigma, \sigma' \in \Sigma(x)$ and $\sigma > \sigma'$

$$\Gamma_x(\sigma, \lambda) - \Gamma_x(\sigma', \lambda)$$

is non-increasing in λ. Since $\Sigma(x) \subset \Re$ this is then a submodular game, in the language of [33]. As a technical constraint, in order to guarantee the existence of best replies, we require that $\Gamma_x(\sigma, \lambda(s))$ be upper semi-continuous in σ.

Note that when $\Gamma_x(\cdot, \cdot)$ is twice differentiable then decreasing differences is equivalent to requiring that $\frac{\partial^2 \Gamma_x(\sigma,\lambda)}{\partial\sigma\partial\lambda_i} \leq 0$ for all $x \in X$ and $i \in [1, m]$. This is a simple extension of Topkis' basic Theorem [45] relating submodularity and cross partial derivatives.

The following notation will be convenient. Define the minimum of $S \in \hat{\Xi}(\Sigma)$ by $[\min(S)](x) = \min[S(x)]$ and $\max(S)$ similarly. Let $\lambda_{\min}(S) = \lambda(\min(S))$ and $\lambda_{\max}(S)$ similarly. Let $\lambda(S)$ be defined to be the set of externality vectors arising from $\lambda(s)$ for strategies arising in S. By Lyapunov's theorem (see e.g.,[21]), $\lambda(S) = [\lambda_{\min}(S), \lambda_{\max}(S)]$ which is a sublattice of Λ.

As we will see in the following examples, non-increasing differences captures the essential aspects of negative externalities.

4 Examples

4.1 Data Networks

An important problem in computer networks arises when many users share a set of common network links.[10] We first consider the simplest case of a single link.

In this case there is one class of customers, $m = 1$. Each player $x \in X$ wants to send

[10] See, e.g., [23][41, 42] for game theoretic analyses of these models.

data across a network link and has utility $\Gamma_x(\sigma, d)$ where $\sigma \in A(x)$ where σ is the (Poisson) rate at which player x sends information. Two common choices for $A(x)$ are $A(x) = [0, \sigma_{\max}]$ and $A(x) = \{0, \sigma_{\max}\}$; in the latter case players are deciding whether or not to use the system. Thus given some strategy vector of transmission rates, s, the externality, $\lambda(s)$ is the total data rate across the link. Because of congestion on the link there is an average (per unit time) total (over packets) delay for all the players – the time between transmission and delivery which depends both on the time taken to traverse the link, and the time spent waiting to use the link – given by $\sigma D^\kappa(\lambda)$, where κ is the capacity of the link (κ^{-1} is the average time it takes for a message to traverse the link) and $D^\kappa(\lambda)$ is the expected waiting time per packet. For example when service times have expectation κ and variance ν^2 we get the M/G/1 delay function

$$D^\kappa(\lambda) = \frac{\lambda(\kappa^2 + \nu^2)}{2(1 - \lambda/\kappa)} + \kappa,$$

which reduces to $D^\kappa(\lambda) = (\kappa - \lambda)^{-1}$ for exponential service times (the M/M/1 queue).

Players have preferences which are increasing in transmission rate and decreasing in delay. For example, one common measure [44] is linear delay costs $c > 0$ which leads to $\Gamma_x(\sigma, \lambda) = \nu_x(\sigma) - c \cdot \sigma Q_x(\lambda)$ which has decreasing differences as long as $D^\kappa(\lambda)$ is nondecreasing in λ, i.e. as long as the queue length is nondecreasing in the arrival rate, which is true for most reasonable queues. A straightforward generalization to nonlinear delay costs $\Gamma_x(\sigma, \lambda) = \nu_x(\sigma) - c(\sigma Q_x(\lambda))$, where $c(\cdot)$ is nondecreasing, also leads to an NAEG. Nonseparable payoff functions of the form $\Gamma_x(\sigma, \lambda) = U_x(\sigma, \sigma Q_x(\lambda))$ are also NAEGs subject to reasonable restrictions.

More general versions of these games occur when there are many links in a network and packets must traverse several links in order to reach their destination. Let V be a set of vertices and L be a set of links. We assume that each player sends packets from her source vertex to her destination vertex via a path, which is a set of links. We first consider the simple case in which paths are predetermined and fixed. In steady state this is essentially a network using the asynchronous transfer mode (ATM) protocol.[11] In this case we let m be the number of paths in use and refer to them by $P(i)$ where $i \in [1, m]$. Thus, X_i is the set of players using path i. Then the total transmission rate on link l is $\lambda_l = \sum_{\{i \ | l \in P(i)\}} \lambda_i$. If κ_l is the capacity of link l then $D^{\kappa_l}(\lambda_l)$ the expected per packet delay at link l and the total expected delay for player $x \in X_i$ transmitting at rate σ is $D^{\kappa}_x(\lambda) = \sum_{l \in P(i)} D^{\kappa_l}(\lambda_l)$. For nondecreasing delay costs $C(\cdot)$, the payoff function is given by $\Gamma_x(\sigma, \lambda) = \nu_x(\sigma) - c(\sigma D^{\kappa}_x(\sigma, \lambda))$ has decreasing differences as long as $D^{\kappa_l}(\lambda_l)$ is nondecreasing in λ_l.

While ATM connections are in use currently on the Internet, most connections currently made are Internet protocol (IP) connections. In this case let m be the number of source-destination pairs and X_i be the set of players transmitting between the pair denoted by i. Under IP, paths are allocated dynamically in response to congestion conditions throughout the net in an attempt to minimize delays. Nonetheless, we can still define $D_x(\lambda)$ to be the sum of the delays along the path used by player x. Once again the game will have decreasing differences and therefore be an NAEG if $D^{\kappa}_x(\lambda)$ is nondecreasing in λ, a typical occurrence.

The current situation on the Internet which is a combination of ATM and IP connections can still be modeled as an NAEG under reasonable assumptions, with obvious extensions of

[11]This and other Internet protocols are discussed in detail in [28].

notation.

4.2 Web Sites and Shared Databases

The games that arise when multiple users access information on a web site or shared database are quite similar in structure to those discussed in the previous section. In this case we let σ represent the number (or complexity) of requests of type i and $D_i(\lambda)$ be the expected delay of a request of type i given the vector λ which summarizes the total number of requests of each type. Then for $x \in X_i$ we have $\Gamma_x(\sigma, \lambda) = U_x(\sigma, D_i(\lambda))$ which (assuming differentiability) has decreasing differences if $\frac{\partial^2 U_x(\sigma,D)}{\partial\sigma\partial D} \leq 0$ and $D_i(\lambda)$ is nondecreasing in λ.

One specific example is when a database is accessed by two classes of users and serves the low priority users ($i = 1$) only when no high priority ($i = 2$) requests are waiting. When this service is nonpreemptive (i.e. once begun requests cannot be interrupted) and $A(x) = \{0, 1\}$ the delay functions are $D_i(0, \lambda) = 0$, $D_1(1, \lambda) = \kappa(\lambda_1 + \lambda_2)/(2[\kappa - \lambda_1][\kappa - \lambda_1 - \lambda_2])$, and $D_2(1, \lambda) = (\lambda_1 + \lambda_2)/(2[\kappa - \lambda_1])$ ([48]). Similar formulae are known for preemptive systems and other types of priority orderings.

4.3 Dialing in to an ISP

As any person dialing in to Internet Service Provider (ISP) or almost any large university has learned, there can be long delays accessing shared computer networks. We first consider a simple (and currently unrealistic) model and then two more realistic versions.

Consider a computer services provider which has k modems. Assume there is one class of users, $m = 1$ and consider a user who must decide whether to dial in from home, or go into her office and use the computer there; thus $S(x) = \{0, 1\}$. If she dials in from home and there are no modems available, then she gets placed in the modem queue. Assume that the distribution of lengths of time that people stay logged on is common knowledge and independent between users. Then, for a given λ the expected wait time for a modem is $D(\lambda)$. This incorporates the case that there are one or more free modems, in which case the wait is zero. For example, if usage times are exponentially distributed then $D(\lambda)$ is well known (but quite complicated), as this is the M/M/k queue with first in first out (FIFO) queuing (see e.g. [22] for the formula); more generally we need to solve for the M/GI/k queuing delay. Let $C_x(\cdot)$ be the player's (nondecreasing) waiting cost and b_x be the cost of driving to work, then we get an NAEG $\Gamma_x(0, \lambda) = -b_x$ and $\Gamma_x(1, \lambda) = -C_x(D(\lambda))$, since $D(\lambda)$ is nondecreasing in λ.

Often there can be multiple classes depending on modem speeds – there may be several different types in the modem pool and certain players have restrictions over the set of modems which they can use. In this case we get an NAEG with several classes.

Now consider a more common scenario in which players who dial in when all the modems are in use get a busy signal. If we assume that people who receive a busy signal immediately redial and continue to do so until connected (many programs do this automatically) and the redial time is short enough, we get essentially the M/M/k queue with random queuing which leads to the same delay function.

Another possibility is that after receiving a busy signal a player immediately exits and does not connect to the system later (that day). In this case the payoff function is $\Gamma_x(0, \lambda) = 0$

and $\Gamma_x(1, \lambda) = p(\lambda)\ v_x - c$ where v_x is the value of connecting, c is the cost of the phone call including delay cost, and $1 - p(\lambda)$ is the probability of getting a busy signal, which is a well known function, but complicated analytically [22].

4.4 Cost Sharing and Market Games

Although in this paper we are focusing on games arising on the Internet, there are many other common examples of NAEGs in Economics. For example, many cost allocation problems induce NAEGs. These include fishermen fishing from a common lake, farmers polluting common water supplies, and various versions of the tragedy of the commons.

We now present a slightly less obvious example of an NAEG which arises when there are a large number of producers producing objects which are identical or close substitutes. Assume that there are m products and let X_i be the set of firms producing product i. Let $A(x)$ be the set of potential production plans by firm x and let $C_x(\sigma)$ be the cost of producing $\sigma \in A(x)$.

If the products are close substitutes and normal, then the market price of good i given the total production vector λ, where λ_i is the total amount of good i produced, is given by $P_i(\lambda)$ which is nonincreasing in λ. In this case the payoffs are given by $\Gamma_x(\sigma, \lambda) = P_i(\lambda)\sigma - C_x(\sigma)$ for $x \in X_i$. Note that this function has nonincreasing differences irrespective of the functional form of $C_x(\cdot)$.

5 Equilibria and Best Reply Dynamics

In this section we analyze the structure of the best reply dynamics for an NAEG. This will allow us to characterize the set of equilibria. Then, using the monotonicity properties of the game, we will examine various ways to reduce the set-valued BR-dynamics to a tractable point valued process, which will allow us to compute (set-valued) convergence explicitly. Although, we do not claim that this is a reasonable model of learning, the convergence of the best reply dynamics will play a key role in understanding the convergence of more general learners as will be discussed in Sections 6 and 7. We begin by analyzing the best reply mapping.

Note that in computing her best replies against strategy profile $s \in \Xi$, all that a player needs to know is $\lambda(s)$. Therefore we define the best reply mapping in terms of λ.

Definition 1 *Let $BR_x(\lambda(s))$ be the set of best replies for player $x \in X$ against strategy profile s, and $BR_x(\lambda(S)) = \bigcup_{s\in\Xi(S)} BR_x(\lambda(s))$.*

We use a standard set ordering for sets of best replies.

Definition 2 *Given $A, B \subseteq \Re$ we say that $A \geq B$ (in the strong set order [47]) if for all $\sigma \in A$, $\sigma' \in B$ then $\max(\sigma, \sigma') \in A$ and $\min(\sigma, \sigma') \in B$.*

For example if A, B in the above definition are single points, then this corresponds to the standard ordering.

Theorem 1 *The set $BR_x(\lambda(s))$ is nonempty, compact and nonincreasing in λ, in the strong set ordering.*

Proof: This proof follows immediately from the upper semicontinuity and decreasing differences of $\Gamma_x(\cdot,\cdot)$. (See [33].) $\square$

Let $BR(\lambda(S))$ be the set-valued mapping from $\hat{\Xi}$ to $\hat{\Xi}$. Since the players' payoffs only depend on $\lambda(S)$ we can write $T(\lambda(S)) = \lambda(BR(S))$ and view the BR dynamics as a set-valued dynamics from Λ to itself.

Corollary 1 *For all $\lambda \in \Lambda$, $T(\lambda)$ is a non-increasing compact interval, e.g.* $[\lambda_{\min}, \lambda_{\max}]$.

Proof: This follows immediately from the previous theorem and Lyapunov's theorem [21]. $\square$

5.1 Equilibria

Although the structure of the equilibria for NAEGs will not be needed for our main analysis, we briefly describe it in this subsection as it is of independent technical interest. In particular, it highlights the differences between NAEGs and supermodular games [46, 33].

Using the properties of $T(\Lambda)$ it is straightforward to analyze the set of equilibria for an NAEG.

Definition 3 *A strategy profile $s \in \Xi$ is a Nash Equilibrium in pure strategies if for almost all x and all $\sigma \in \Sigma(x)$, $\Gamma_x(s(x), \lambda(s)) \geq \Gamma_x(\sigma, \lambda(s))$. Let NE be the set of Nash equilibria.*[12]

First we show that we can reduce the study of the set NE to that of $\lambda(NE)$ which are the externality vectors associated with the set of equilibria.

Theorem 2 *If $\lambda \in T(\lambda)$ then there exists a strategy profile $s \in NE$ with $\lambda(s) = \lambda$.*

[12] As is well known, (e.g. [40]) it is not necessary to consider mixed strategies.

Proof: Let $S = BR(\lambda)$, then by assumption $\lambda \in \lambda(S)$ and therefore there is an $s \in \hat{\Xi}(S)$ such that $\lambda(s) = \lambda$. Since $s \in \Xi(S)$ it is a best reply to itself and therefore in NE. $\square$

Now we show that the structure of NE is quite different from the structure of the Nash equilibria in supermodular games.

Definition 4 *A set $A \subset \Lambda$ is a 'necklace' if for all $\lambda, \lambda' \in A$ such that $\lambda \leq \lambda'$, then for any $\lambda'' \in [\lambda, \lambda']$, $\lambda'' \in A$.*

A necklace can be viewed as a string of gems (intervals), such that no two points from different gems are ordered.

Theorem 3 *$\lambda(NE)$ is a necklace.*

Proof: Assume that $\lambda, \lambda' \in \lambda(NE)$ and that $\lambda < \lambda'$. Choose any $\lambda'' \in [\lambda, \lambda']$. Now $\lambda \in T(\lambda)$ and $\lambda' \in T(\lambda')$. Since $T(\cdot)$ is decreasing and consists of intervals, this implies that $\lambda'' \in T(\lambda')$. $\square$

Lastly we note that an NAEG has a Nash equilibria if the payoff function is continuous.

Theorem 4 *If $\Gamma_x(\cdot, \lambda)$ is continuous then $NE \neq \emptyset$.*

Proof: This is proven in [40] for general non-atomic games. In our case it follows immediately, from the monotonicity of T. $\square$

In contrast to supermodular games, discontinuous NAEGs need not have an equilibrium as Tarski's theorem only implies that the best reply mapping has a 2-cycle, which may not be and equilibrium.

5.2 Best Reply Convergence

In this subsection we analyze the convergence of best reply dynamics since, as we will see in the next two sections, such convergence implies that for NAEGs more general learning procedures will also converge.

Since $BR(S)$ is a monotone mapping $S^* = lim_{n \to \infty} BR^n(\sigma)$ is well defined, where $BR^{n+1}(S) = BR(BR^n(S))$ and $BR^0(S) = S$.

Definition 5 *An Externality Game is BR-stable if* $\forall x \in X, |S^*(x)| = 1$ *a.e., i.e. the limit set consists of a single strategy.*

However, direct computation of set-valued convergence can be quite difficult. We now reduce this problem to one involving only point valued functions. Let $T_-(\lambda) = \min[T(\lambda)]$ and $T_+(\lambda) = \max[T(\lambda)]$ and note that $T(\lambda) = [T_-(\lambda), T_+(\lambda)]$. Define $T_\pm(\lambda) = T_+(T_-(\lambda)$ and $T_\mp(\lambda) = T_-(T_+(\lambda)$. We can compute BR-stability directly from the (point valued) mappings $T_\pm$ and $T_\mp$.

Theorem 5 *An Externality Game is BR-stable if and only if* $\lim_{n \to \infty} T^n_\pm(\Lambda) = \lim_{n \to \infty} T^n_\mp(\Lambda)$.

Proof: This follows from the proof of Theorem 3 where it is shown that $T(\Lambda) = [T_-(\Lambda), T_+(\Lambda)]$. Applying T repeatedly to both sides of this equation yields $T^{2n}(\Lambda) = [T^n_\mp(\Lambda), T^n_\pm(\Lambda)]$. Taking the limit as $n \to \infty$ completes the proof. □

Thus we have reduced the problem of set valued convergence to one with single valued functions. We now reduce the problem further to the study of two specific trajectories.

Theorem 6 *If* $\lim_{n\to\infty} T_{\pm}^{n}(0) = \lim_{n\to\infty} T_{\mp}^{n}(N)$ *then the game is BR-stable. (Note that the limits exist by monotonicity.)*

Proof: This follows immediately since $T_{\pm}(\min[A]) \leq T_{\pm}(A)$ and similarly for $T_{\mp}$. □

When $T(\cdot)$ is single valued then the analysis is particularly simple.

Corollary 2 *Assume that* $|T(\lambda)| = 1$ *for all* λ*, then if* $\lim_{n\to\infty} T^{n}(0)$ *exists the game is BR-stable.*

Proof: This follows immediately from the previous theorem. □

Thus it is easy to compute BR-stability of an NAEG. For example, in a symmetric Externality Game, since O is symmetric, $T_{\pm}, T_{\mp}$ are also symmetric, and by defining $T_{\pm}(x) = T_{\pm}(x, x, ..., x)$ (resp. $T_{\mp}$) the problem of computing stability is reduced to that of a 1-dimensional monotone mapping, which is trivial to analyze. See [12, 10] for examples.

6 D-Solvability: Learning with Observability

Our goal in this section is to relate BR-stability to D-solvability to provide an easy test for whether a game is learnable. This is important for understanding learning behavior in a repeated game, since as discussed in Section 2 Milgrom and Roberts [33] have shown that in settings in which players can observe other players' strategies, many learning procedures converge to the serially undominated set (defined below) which implies that convergence is unique if the game is D-solvable.

Define the operator $D_x(S)$ to be the set of undominated strategies (with respect to pure strategies) for player x against strategy set S.

Definition 6 *For player x, (pure) strategy σ dominates (pure) strategy σ' relative to $S \in \hat{\Xi}$ if and only if $\Gamma_x(\sigma, \lambda) > \Gamma_x(\sigma', \lambda)$ for all $\lambda \in \Lambda(S)$. The set of strategies which are not dominated are denoted $D_x(S)$.*

Inductively define $D^1(S) = D(S)$ and $D^k(S) = D(D^{k-1}(S))$. Since D is a monotonic set operator $D^\infty(\Sigma)$ is well defined. (For details see [34].) For adaptive learners the dynamic play essentially follows this repeated elimination of dominated strategies.

Definition 7 *A game is D-solvable if $|D^\infty(\Sigma)(x)| = 1$ for almost all x.*

D-Solvability also implies that a game has very strong equilibrium properties. For example a D-solvable game has a unique rationalizable strategy.

It is well known that D-solvability implies BR-stability. However, the converse is untrue in general. (See e.g. [35].) Interestingly, in NAEGs the converse is true.

Theorem 7 *An Externality Game is BR-stable if and only if it is D-solvable.*

Proof: (This argument is similar to that of [34] for supermodular games.) The if direction follows immediately since a best reply can not be dominated. To prove the other direction we will show that any undominated strategy must lie between the extreme best replies, and thus $\lambda(D(S)) \subseteq \lambda(BR(S))$.

Assume that there exists a strategy σ which is not a best reply, but is undominated, with respect to the set S. Let σ_- be the smallest member of $BR_x(S)$ and σ_+ the largest. Assume that $\sigma \notin [\sigma_-, \sigma_+]$. For convenience assume that $\sigma < \sigma_-$. (The opposite case is treated similarly.) Since the best reply mapping is non-increasing in λ it follows that $\sigma_- \in BR_x(\sigma_+, \lambda_{\max}(S))$ and $\sigma \notin BR_x(\sigma_+, \lambda_{\max}(S))$ thus

$$\Gamma_x(\sigma_-, \lambda_{\max}(S)) - \Gamma(\sigma, \lambda_{\max}(S)) = \Delta > 0.$$

However by decreasing differences we must also have

$$\Gamma_x(\sigma_-, \lambda_{\min}(S)) - \Gamma(\sigma, \lambda_{\min}(S)) > 0,$$

and therefore σ is dominated by σ_-, a contradiction. □

Thus, one can easily test the D-solvability of an NAEG by applying Theorem 5 or Corollary 1 thus showing convergence of a wide variety of synchronous learning procedures, as discussed in Section 2. Section 7 demonstrates uses of this result.

7 O-Solvability: Decentralized Learning

While D-solvability is an extremely strong property, Friedman and Shenker [14, 15] argue that it does not guarantee convergence of learners in decentralized settings such as the Internet. Thus, we are led to consider the question of when NAEGs are learnable in these settings. In [15] it is shown that even in noisy asynchronous decentralized settings with little common knowledge, 'reasonable' learners will converge to the serially unoverwhelmed set (defined below).

In this section we show that for a large class of NAEGs BR-stability is equivalent to O-solvability, which implies that the serially unoverwhelmed set is a singleton and thus reasonable learners will converge to the (unique) Nash equilibrium. First we define our solution concept.

Definition 8 *For player x, strategy σ overwhelms strategy σ' relative to $S \in \hat{\Xi}$ if and only if $\Gamma_x(\sigma, \lambda) > \Gamma_x(\sigma', \lambda')$ for all $\lambda, \lambda' \in \Lambda(S)$. The set of strategies which are not overwhelmed are denoted $O_x(S)$.*

Define $O^k(S) = O(O^{k-1}(S))$ with $O^1(S) = O(S)$, and note that $O^\infty(\Sigma) = \lim_{k\to\infty} O^k(\sigma)$ is well defined due to the monotonicity of $O(\cdot)$.

Definition 9 *A game is O-solvable if $|O^\infty(\Sigma)(x)| = 1$ for almost all x.*

O-solvability is a very strong condition. In the learning context it implies that even in an asynchronous game where a player is unaware of her opponents strategies, she can still tell which strategies are dominated. This allows for the convergence of very simple learning algorithms that may use very little information.

Now we formally define the class of simple NAEGs. In general these encompass 'participation' games where players have only 2 choices, whether or not to participate. Several examples of these were given in Section 3.

Definition 10 *An NAEG is simple if for all $x \in X$, $A(x) = \{\sigma_-, \sigma_+\}$ and $\Gamma_x(\sigma_-, \lambda)$ is nondecreasing in λ, where $\sigma_- < \sigma_+$.*

The typical example of a simple NAEG arises when the payoff for non participation is 0, i.e. $\Gamma_x(\sigma_-, \lambda) = 0$. In the general case $\Gamma_x(\sigma_-, \lambda)$ is nondecreasing in λ which implies that the value of nonparticipation is nondecreasing in the congestion. This may arise when nonparticipation implies using the default resource, which is less congested when more players are using the alternate resource. Section 7 provides an example of this.

Our main result for simple NAEGs connects their O-solvability to D-solvability and thus by the previous theorems their BR-stability.

Theorem 8 *If a Simple NAEG is D-solvable, then it is O-solvable.*

Proof: First note that by definition $D_x(S) \subseteq O_x(S)$. Assume that σ dominates σ' with respect to S and that $\sigma > \sigma'$ (the other case follows similarly). Thus $\Gamma_x(\sigma, \lambda) - \Gamma_x(\sigma', \lambda) > 0$ for all $\lambda \in \lambda(S)$. Consider the pair $\lambda > \lambda' \in \lambda(S)$. Then $\Gamma_x(\sigma, \lambda') - \Gamma_x(\sigma', \lambda') > 0$ and by decreasing differences we see that $\Gamma_x(\sigma, \lambda) - \Gamma_x(\sigma', \lambda') > 0$. When $\lambda < \lambda'$ we get $\Gamma_x(\sigma, \lambda') - \Gamma_x(\sigma', \lambda) > 0$ since $\Gamma_x(\sigma_{\min}, \lambda)$ is nondecreasing in λ. Thus $D_x(S) = O_x(S)$ and iterating this completes the proof. $\square$

Thus, we can easily check whether a simple NAEG is O-solvable by testing its BR-stability.

Corollary 3 *If a Simple NAEG is BR-stable, then it is O-solvable.*

8 Applications

Consider the following simple (ATM) printing network with three links a, b, c, which are joined as the letter Y and two paths which are used to send information. Players at location $i = 1$ send packets across their local link a to the common printer, which we model as link b. Players at location $i = 2$ send packets across link c to the same printer b. For simplicity assume that each player chooses to print locally $\sigma = 0$ or via the network on the high quality printer $\sigma = 1$. Let $X_1 = X_2 = [0, 1]$ with associated measures μ_i which are the Lebesgue measure. Assuming exponential service times leads to $D_1(\lambda) = (\mu_a - \lambda_1)^{-1} + (\mu_c - \lambda_1 - \lambda_2)^{-1}$

and $D_2(\lambda) = (\mu_c - \lambda_2)^{-1} + (\mu_c - \lambda_1 - \lambda_2)^{-1}$. Assuming that the players value of using the high quality printer is $10x$ and that they have linear delay costs the payoff functions are $\Gamma_x(0, \lambda) = -c$ and $\Gamma_x(1, \lambda) = 10x - cD_i(\lambda)$ for $x \in X_i$, where we assume that the local printer has an expected printing time of 1.

It is easy to solve for the best reply mapping by finding the indifferent player, who will satisfy $x_i = c(D_i(\lambda) - 1)/10$. Thus,

$$T(\lambda) = (1 - c(D_1(\lambda) - 1)/10, \quad 1 - c(D_2(\lambda) - 1)/10).$$

Recalling the delay formulae from the previous paragraph, one notices that these formulae are quite complex algebraically and thus without using the structure of NAEGs the analysis of equilibria and solvability is quite difficult. However, using Corollaries 3 and 2, it is easy to check for O-solvability for specific values of the capacities with a standard commercial spreadsheet.

For example, when $\mu_a = 2, \mu_b = 3, \mu_c = 4$ and $c = 1$ the trajectory starting at $(0, 0)$ converges very rapidly to approximately $(0.835676263, 0.889414738)$, showing both that the equilibrium is unique and that the game is O-solvable, and thus easily learnable even in asynchronous settings.[13] Whereas for $\mu_a = \mu_b = \mu_c = 1.9$ the dynamics starting at $(0, 0)$ alternates between $(0, 0)$ and $(0.894736842, 0.894736842)$ showing that the game is not learnable. However for $\mu_a = \mu_b = \mu_c = 2$ the dynamics converges to $(0.733457526, 0.733457526)$ demonstrating that the dividing line between unlearnable and learnable can be quite abrupt.

[13]The numerical results converge to 3 decimal places in only 4 iterations and 7 decimal places in 10 iterations. Although the numerical results alone can not prove convergence, a simple bound on the derivative of the mapping can be used to rigorously prove convergence with numerical computations.

In fact, it is not coincidental that the game is learnable for large μ and unlearnable for small μ as the following theorem demonstrates.

Theorem 9 *For any data network game as defined in Section 3.1 the exists a vector $\overline{\mu}$ such that for any $\mu > \overline{\mu}$ the game is BR-stable and thus D-solvable. If in addition the game is simple, then it is also O-solvable.*

Proof: One can prove this by showing that for sufficiently large μ the best reply mapping $T(\lambda)$ is a contraction mapping. This is a straightforward extension of Theorem 3 in [12]. □

Finally, note that these results are unchanged even if there is congestion at the local printers. This is because the externality at local printer i is given by $1 - \lambda_i$ and thus the payoff for $\sigma = 0$ is increasing in λ.

Thus, if a network has sufficient capacity it will be learnable by reasonable players without much difficulty.

9 Conclusions/Extensions

We have provided a simple and computable test for the learnability of a large class of non-atomic games arising in the context of the Internet. Similar results arise in finite player games, although the analytic computations are (apparently) somewhat more complex. Thus, the basic intuitions in this paper should transfer to the finite player case.

Also, in games with a finite number of players, the network protocol becomes an important tool in designing learnable games. For example, the "fair share" service discipline [42] induces games which are learnable (O-solvable) for a wide variety of Internet games

[15]. Thus, it seems possible that one can design networks which are intrinsically learnable, thereby reducing the need for applications to use sophisticated optimization algorithms by simplifying the process of attaining equilibria with its many attendant benefits.[14]

References

[1] R. J. Aumann and L. S. Shapley. *Values of Non-Atomic Games.* Princeton University Press, Princeton, N.J., 1974.

[2] Y. Chen. Asynchronicity and learning in cost sharing mechanisms. Mimeo, 1996.

[3] David D. Clark, Scott Shenker, and Lixia Zhang. Supporting real-time applications in an integrated services packet network: architecture and mechanism. In *Proceedings of Sigcomm 92*, pages 14–26, Baltimore, Maryland, August 1992. ACM. *Computer Communication Review*, Volume 22, Number 4.

[4] Ron Cocchi, Scott Shenker, Deborah Estrin, and Lixia Zhang. Pricing in computer networks: Motivation, formulation, and example. *Transactions on Networking*, 1(6), December 1993.

[5] Alan Demers, Srinivasan Keshav, and Scott Shenker. Analysis and simulation of a fair queueing algorithm. *Journal of Internetworking*, 1(1):3–26, January 1990.

[14] These include the reduction of computation for optimizing resource usage, and the ability to analyze the social welfare properties of a network configuration.

[6] Donald F. Ferguson. *The Application of Microeconomics to the design of resource allocation and control algorithms in Distributed Systems.* PhD thesis, Dept. of Computer Science, Columbia University, New York, New York, 1989.

[7] Donald F. Ferguson, Christos Nikolaou, and Yechiam Yemini. An economy for flow control in computer networks. In *Infocom*, pages 110–118, Ottawa, Canada, April 1989. IEEE.

[8] D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. Mimeo, 1996.

[9] D. Foster and P. Young. Learning with hazy beliefs. Mimeo, 1996.

[10] E. J. Friedman. Dynamics and rationality in ordered externality games. *Games and Economic Behavior*, 16:65–76, 1996.

[11] E. J. Friedman, Amy Greenwald, and S. Shenker. Responsive learning in network games. mimeo, 1998.

[12] E. J. Friedman and A. S. Landsberg. Short run dynamics of multi-class queues. *Operations Research Letters*, 14:221–229, 1993.

[13] E. J. Friedman and A. S. Landsberg. Long run dynamics of queues. *Operations Research Letters*, 18:185–191, 1996.

[14] E. J. Friedman and S. Shenker. Synchronous and asynchronous learning by responsive learning automata. mimeo., 1995.

[15] E. J. Friedman and S. Shenker. Learning and implementation in the internet. mimeo, 1997.

[16] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.

[17] A. Gupta, D. O. Stahl, and A. B. Whinston. Managing the Internet as an economic system. Technical report, Department of Economics, University of Austin at Texas, July 1994.

[18] M. W. Hirsch. *Differential Topology*. Springer-Verlag, New York, 1976.

[19] M. T. Hsiao and A. Lazar. A game theoretic aproach to dezentralized flow control of marko vian queueing networks. In Courtois and Latouche, editors, *Performance'87*, pages 55–73. North-Holland, 1988.

[20] M. Kandori, G. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61:29–56, 1993.

[21] J.L. Kelley. *General Topology*. Van Nostrand, New York, 1955.

[22] L. Kleinrock. *Queueing Systems, Vol I*. John Wiley, New York, 1975.

[23] Y. Korilis, A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE Journal on Selected Areas in Communications*, 13(7), 1995.

[24] Y. A. Korilis and A. A. Lazar. On the existence of equilibria in noncooperative optimal flow control. CTR Technical Report 340-93-13, Center for Telecommunications Research, Columbia University, New York, April 1993.

[25] Y. A. Korilis, A. A. Lazar, and A. Orda. The designer's perspective to noncooperative networks. In *Infocom*, Boston, Massachusetts, April 1995.

[26] Y. A. Korilis, A. A. Lazar, and A. Orda. The role of the manager in a noncooperative network. In *Infocom*, San Fransisco, California, March 1996.

[27] R. Lagunoff and A. Matsui. An "anti-folk theorem" for a class of asynchronously repeated coodination games. Mimeo, 1995.

[28] J. Mackie-Mason and H. Varian. Economic faqs about the Internet. *Journal of Economic Perspectives*, 8(4):75–96, 1994.

[29] Jeffery Mackie-Mason and Hal R. Varian. Pricing congestible network resources. Technical report, University of Michigan, Michigan, USA, July 1994.

[30] Jeffrey K. MacKie-Mason and Hal Varian. Pricing the internet. In Brian Kahin and James Keller, editors, *Public Access to the Internet.* ACM, Boston, Massachusetts, May 1993. version of February, 1994.

[31] Steven McCanne, Van Jacobson, and Martin Vetterli. Receiver-driven layered multicast. In *Proceedings of Sigcomm 96*, Palo Alto, California, August 1996.

[32] H. Mendelson and S. Whang. Optimal incentive-compatible priority pricing for the m/m/1 queue. *Operations Research*, 38(5):870–883, September–October 1990.

[33] P. Milgrom and J. Roberts. Rationalizability, learning and equilibrium in games with strategic complementarities. *Econometrica*, 58:1255–1278, 1990.

[34] P. Milgrom and J. Roberts. Adaptive and sophisticated learning in repeated normal form games. *Games and Economic Behavior*, 3:82–100, 1991.

[35] H. Moulin. Dominance solvability and Cournot stability. *Mathematical Social Sciences*, 7:83–102, 1984.

[36] John Murphy and Liam Murphy. Bandwidth allocation by pricing in ATM networks. Technical report, Dublin City University, Ireland, July 1994.

[37] J. Nagle. On packet switches with infinite storage. RFC 970, Internet Engineering Task Force, December 1985.

[38] Y. Rekhter, P. Resnick, and S. Bellovin. Financial incentives for route aggregation and efficient address utilization in the internet. preprint, ATT, 1996.

[39] B. A. Sanders. An incentive compatible flow control algorithm for rate allocationin computer networks. *IEEE Transactions on Computers*, C-37(9):1067–1072, 1988.

[40] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7:295–300, 1973.

[41] S. Shenker. Efficient network allocations with selfish users. In P. J. B. King, I. Mitrani, and R. J. Pooley, editors, *Performance '90*, pages 279–285. North-Holland, New York, 1990.

[42] S. Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3:819–831, 1995.

[43] J. Maynard Smith. *Evolution and the theory of games.* Cambridge University Press, Cambridge, 1982.

[44] S. Stidham. Pricing and capacity decisions for a service facility: Stability and multiple local optima. *Management Science*, 38(8):1121–1139, 1992.

[45] D.M. Topkis. Minimizing a submodular function on a lattice. *Operations Research*, 26:305–321, 1978.

[46] D.M. Topkis. Equilibrium points in non-zero sum n-person submodular games. *SIAM J. Control and Optimization*, 17(6):773–787, 1979.

[47] A. Veinott. Lattice programming. unpublished notes, 1989.

[48] R.W. Wolff. *Stochastic Modelling and the Theory of Queues.* Prentice-Hall international series in industrial and systems engineering. Prentice Hall, Inc., Englewood Cliffs, New Jersey, 1989.