Discussion Paper No. 1165

DEVELOPING SYMBOLIC DYNAMICS TO MEASURE THE COMPLEXITY OF REPEATED GAME STRATEGIES BY TOPOLOGICAL ENTROPY

ERIK M. BOLLT AND MICHAEL A. JONES

September 1996

DEVELOPING SYMBOLIC DYNAMICS TO MEASURE THE COMPLEXITY OF REPEATED GAME STRATEGIES BY TOPOLOGICAL ENTROPY*

ERIK M. BOLLT AND MICHAEL A. JONES

ABSTRACT. State spaces, directed graphs, and transition matrices are implemented to consider sequences of play of repeated games as a dynamical system on the symbol space, defined as the outcomes of the stage game. Strategies are described as subshifts of this symbol space. The complexity of of bounded recall strategies is equated with the topological entropy of an associated transition matrix. The topological entropy measures how complicated a strategy is by what paths of the extensive form game are possible under the strategy. This measure is used to compare the complexity of different strategies.

1 INTRODUCTION

The complexity of finite automata has been applied to measure the complexity of the strategies which they mimic. Specifically, the number of states of a minimal finite automaton that plays a strategy is defined to be the complexity of that strategy (*e.g.* Kalai and Stanford, 1988). This approach has led in a number of different directions. One of the motivating factors for this research has been the intuition that highly complex strategies may not be frequently used by players. For instance, complexity costs may deter individuals from using different strategies. Also, the number of states of a finite automaton correspond with the minimal amount of information or memory needed to implement the strategy.

Neyman (1985) and Rubinstein (1986) determine what solutions are possible for repeated Prisoner's Dilemma games when there are restrictions on the complexity of finite automata. Lehrer (1988) examines what equilibrium payoffs are possible under stationary bounded recall strategies in repeated games when players have different memory capacities. More recent work includes Johnson (1995) which introduces a finer measure of complexity associated with the algebraic complexity of semigroups. Other recent work includes defining strategic entropy to measure the complexity of mixed strategies in two person, zero-sum games, where one player is restricted to a smaller set of strategies (Neyman, 1996). A more complete survey of other uses of complexity can be found in Kalai (1990) and (1995).

Date: September 1996.

Key words and phrases. Complexity, Topological Entropy, Symbolic Dynamics, Repeated games.

This paper focuses on a slightly different, but related question associated with repeated game strategies. When strategies are limited to stationary bounded recall strategies, can one determine how complex the possible outcomes of a repeated game can be? For our purposes, the complexity of a strategy or strategy profile is equated with the complexity of the possible outcomes of the repeated game under that strategy or strategy profile.

To analyze the complexity of the possible outcomes under bounded recall strategies, the action space of the stage game defines an alphabet or symbol space. This alphabet leads to interpreting the possible outcomes under a strategy as a subshift of symbol space. By applying ideas from dynamical systems, we measure the complexity of the outcomes by determining the topological entropy of a transition matrix associated with a strategy. To measure the complexity of a strategy profile, an algebra is used to create a transition matrix for the strategy profile from the transition matrices of the individual strategies that make up the profile.

A bounded recall strategy dependent on the last k rounds of play uses less information than a bounded recall strategy dependent on the last $k+1$ rounds. Consequently, their corresponding transition matrices are of different dimensions. To compare matrices of different dimensions, a lift map is defined. This lift represents a matrix as a "higher" dimensional matrix, so that matrices can be operated on to form the transition matrix of a strategy profile.

Symbolic dynamics and associated directed graphs and matrices, comparable to Markov transition matrices, are introduced in Section 2. Section 3 introduces the lifting technique. Section 4 defines an algebra on the transition matrices and utilizes the lift to compare matrices of different dimensions. Section 5 defines topological entropy, its use as a measure of the complexity of a strategy, and proves that this measure of complexity is well-defined, easy to calculate, and a natural way to consider complexity. Different bounded recall strategies are compared in Section 6. Examples throughout the paper introduce the new material and allow for the comparison of strategies of a repeated Prisoner's Dilemma game.

2. Notation and Symbolic Dynamics

Let $N = \{1, 2, \ldots, n\}$ be the set of players. The stage game is defined by the set of actions available to Player i, $D_i = \{d_{i_1}, d_{i_2}, \ldots, d_{i_{m_i}}\}$, at every round. Notice that D_i consists of m_i possible actions. The set $D = D_1 \times D_2 \times \cdots \times D_n$ is the action space of the stage game and the cardinality of D is $m = \prod_{i=1}^{n} m_i$. Normally, a utility function is defined for each player mapping D into the real numbers. We do not require utility functions for the players at this time since we are interested in what the strategies are, not whether they are Nash equilibria, *etc.* We restrict our consideration to players' bounded recall strategies that are independent of time. Lehrer (1988) defines these bounded recall strategies as stationary. The following development assumes that the repeated game lasts an infinite number of rounds. The analysis can easily be truncated for finitely repeated games.

Definition 2.1. A bounded recall strategy of k rounds, for Player i, is a strategy which depends only on the last k actions and is independent on the round of the game. For every round t, the strategy

$$S_i^t = S_i : H^k \longrightarrow \Delta D_i, \tag{2.1}$$

is a bounded recall strategy of k rounds, where H^k is the set of all k length histories and ΔD_i is the set of probability distributions on D_i.

To simplify the initial analysis and presentation, we assume that there is a common knowledge initial state which contains sufficient terms for every players' bounded recall strategy to be applicable. Define the symbol space on D as $\Sigma_D = \{\sigma_0.\sigma_1\sigma_2\ldots|\sigma_i \in D\}$, which is also called the full shift. A point in symbol space is a possible path of play. The space Σ_D represents all possible sequences of play, often denoted by H^∞, the set of all possible histories of the infinitely repeated game. The following definition relates strategies to symbolic dynamics.

Definition 2.2. The Bernoulli shift on Σ_D is a map $s : \Sigma_D \longrightarrow \Sigma_D$ defined by

$$s(\sigma) = s(\sigma_0.\sigma_1\sigma_2\sigma_3\ldots) = \sigma_1.\sigma_2\sigma_3\ldots. \tag{2.2}$$

A strategy for Player i restricts the possible sequences of play. Hence, a strategy can be represented by all the possible paths of play that can occur under that strategy. This limitation on possible paths of play defines a subspace of Σ_D. Another way to consider this subspace is as the branches of the extensive form of the repeated game that can be reached by the strategy profile of the players.

Definition 2.3. A subshift Σ is a subspace of Σ_D which is invariant under the full shift, *i.e.*, if $\sigma \in \Sigma$ then $s(\sigma) \in \Sigma$.

Every bounded recall strategy can be represented by a subshift. This is apparent since the strategy restricts what actions can be played after a certain history has been played. To understand the intuition and to introduce the notion of transition matrices, consider a simple strategy that only depends on the realized actions from the previous round. The strategy for Player i is a function from D to probability distributions on D_i. The set D doubles as the enumerated state space. Player i's strategy can be represented as a directed graph and transition matrix between the possible states. The transition matrix, M, has entries $m_{s,t}$ equal to 0 or 1. The entry $m_{s,t} = 0$ if state t cannot be reached by state s under Player i's strategy. The entry is 1 if state t can be reached by state s under Player i's strategy. The following example, using a simple strategy dependent only on the realized actions of the preceding round, describes the directed graph, transition matrix, and subsequent subshift.

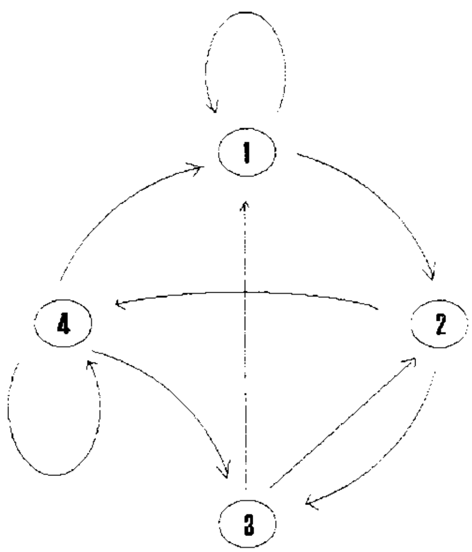

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}$$

FIGURE 1. Player 1's Tit-for-tat Strategy Transition Matrix and Directed Graph

Example 2.1. Tit-for-tat for Player 1 in the infinitely repeated Prisoner's Dilemma

The stage game for the Prisoner's Dilemma is defined by $N = \{1, 2\}$ and $D_1 = D_2 = \{c, d\}$, where c represents 'cooperate' and d represents 'defect.' (Although there are certain requirements on the payoffs for the stage game to be a Prisoner's Dilemma, we suppress the payoffs to emphasize the measurement of the complexity of the strategy, independent of the utility functions.) The tit-for-tat strategy for Player 1 is a map $S_1 : D \to \Delta D_1$ defined by $S_1(c, c) = c$, $S_1(c, d) = d$, $S_1(d, c) = c$, and $S_1(d, d) = d$. The image c is equated with the probability distribution that plays c with certainty, similarly for d.

Consider the set D as a set of possible states. Enumerate these states as $1 = (c, c)$, $2 = (c, d)$, $3 = (d, c)$, and $4 = (d, d)$. Alternatively, we can define the tit-for-tat strategy as a point-to-set correspondence, $\hat{S}_1 : D \to D$, which maps to possible plays in the next round. Therefore, $\hat{S}_1(c, c) = \{(c, c), (c, d)\}$ since Player 1 will play c and Player 2 may cooperate or defect. Equivalently, $\hat{S}_1(c, c) = \{(c, *)\}$, which indicates that Player 2's strategy is unknown, and its outcome could be either c or d. Using the enumerated states, a directed graph and transition matrix between these states is defined (in Figure 1) which represents the tit-for-tat strategy for Player 1 (realize that this tit-for-tat strategy would generate a different directed graph and matrix if used by Player 2).

The transition matrix, A, shows that (in terms of the numbered states) the tit-for-tat strategy for Player 1 restricts possible sequences of play to those which: state 1 is followed by state 1 or 2, state 2 is followed by state 3 or 4, state 3 is followed by state 1 or 2, and state 4 is followed by state 3 or 4. For example, the sequence of play 1.23412341234... cannot occur since state 4 follows state 3. However, the state 1.23123123... is possible and would be part of the subshift that defines the possible sequences of play by Player 1's tit-for-tat strategy.

Proposition 2.1. *A transition matrix A, generated from a bounded recall strategy of a single round, defines all possible transitions between the numbered states of D and, thereby, generates a subshift Σ_A of Σ_D.*

Proof. The subshift Σ_A includes all points σ which represent all possible trajectories through the associated directed graph. $\square$

When a strategy is dependent on more information than the realized action from the last round, the technique from the above example may not give fine enough results. The number of states may be too restrictive, but can be split into more states. As an extension of Example 2.1, let the strategies S_2 and T_2, for Player 1, both depend on the actions from the last two previous rounds. Specifically, let $S_2((c,c),(c,c)) = d$ and $S_2((d,d),(c,c)) = c$. The point-to-set correspondence representation would yield $\hat{S}_2((c,c),(c,c)) = \{(d,*)\}$ and $\hat{S}_2((d,d),(c,c)) = \{(c,*)\}$. The strategy S_2 allows for all states to follow state 1. The strategy T_2 yields similar results if $T_2((c,c),(c,c)) = c$ and $T_2((d,d),(c,c)) = d$. The induced correspondence yields $\hat{T}_2((c,c),(c,c)) = \{(c,*)\}$ and $\hat{T}_2((d,d),(c,c)) = \{(d,*)\}$. Again, any state may follow state 1. However, there is an obvious difference between the two strategies S_2 and T_2. The strategy S_2 allows for sequences of play with the triples 113, 114, 411, and 412. The subshift associated with T_2 contains sequences with the triples 111, 112, 413, and 414. This implies that the states under consideration should really be $D \times D$, not just D. Bounded recall strategies of more than a single round are considered in the next section

3. The Lift

From the previous section, a strategy on the action space D generates a subshift Σ consisting of all the allowed sequences of play adhering to the strategy, and Σ is a subset of Σ_D, the set of *all* permutations of *all* sequences of play. The only strategies treated so far, are according to the transition rules of an $m \times m$ matrix A which defines a single round bounded recall strategy. As described by Proposition 2.1, all *allowed* sequences of play, generated by such a transition matrix, are elements of the subshift Σ_A, a subset of Σ_D. However, bounded recall strategies which rely on more information than just the single previous round cannot be described in this manner.

Nonetheless, bounded recall strategies which rely on the k previous actions also define a subshift, which we denote $\Sigma' \subset \Sigma_D$. However, the grammar of such a multiple-step strategy's subshift cannot be described by any $m \times m$ transition matrix. To adequately describe the grammar for the symbol dynamics of Σ' in terms of an allowed set of transitions, here we develop the *lift* of the symbol space Σ_D to a larger ("higher dimensional") symbol space in which the k-step strategy is generated by an $m^k \times m^k$ transition matrix B. The matrix B defines the strategy as the subshift Σ_B, a subset of the larger symbol space, in a straight-forward manner.

Transition rules between actions in D, that incorporate information about the previous k actions, necessitate a state space larger than D. Strategies which incorporate k previous states are best represented by a Bernoulli shift map on a higher dimensional symbol space. To define a k-step memory dependent transition rule of the m symbols in D requires a transition rule on the m^k symbols of D^k. Denote each k-tuple of symbols from D as a symbol in D^k. Let Σ_{D^k} denote the symbol space of *all* infinite sequences of symbols from D^k. As defined below, the k-step memory dependent grammar strategy on D is represented uniquely by a Bernoulli shift on the subshift $\Sigma_B \in \Sigma_{D^k}$.

The Bernoulli shift map applied to a point $\sigma \in \Sigma_D$, where k states are recorded, may be written as

$$s(\sigma) = s(\underbrace{\sigma_0\sigma_1\sigma_2...\sigma_{k-1}}_{k \text{ symbols}}.\sigma_k...) = \underbrace{\sigma_1\sigma_2\sigma_3...\sigma_{k-1}\sigma_k}_{k \text{ symbols}}.\sigma_{k+1}... \qquad (3.1)$$

Writing the decimal immediately after the k^{th} symbol emphasizes the most recent round of play.

Consider D^k as the set of m^k symbols achieved by labeling all permutations of m states over k consecutive rounds of play. Define the symbol space with labels from D^k as $\Sigma_{D^k} = \{\sigma'_0.\sigma'_1\sigma'_2 \ldots | \sigma'_i \in D^k\}$. Further, define the lift function $l : \Sigma_D \to \Sigma_{D^k}$ by collecting groups of k symbols from D in the following "overlapping" manner where $l(\sigma) = \sigma'$. For the point $\sigma = \sigma_0.\sigma_1\sigma_2 \ldots \sigma_{k-1}\sigma_k\sigma_{k+1} \cdots \in \Sigma_D$, the image of σ under l is $\sigma' = \sigma'_0.\sigma'_1\sigma'_2 \cdots \in \Sigma_{D^k}$ where $\sigma'_0 = \sigma_0\sigma_1\sigma_2...\sigma_{k-1}$, $\sigma'_1 = \sigma_1\sigma_2\sigma_3 \ldots \sigma_k$, $\sigma'_2 = \sigma_2\sigma_3\sigma_4 \ldots \sigma_{k+1}$, *etc.* Notice that each σ'_n is a point in D^k, as required.

Example 3.1. Tit-for-two tat strategy for Player 2

Consider the Player 2 bounded recall strategy of 2 rounds, the tit-for-two tat strategy, where Player 2 follows the strategy: "I do what you did two rounds ago." As in Example 2.1, $N = \{1, 2\}$ and $D = D_1 \times D_2$ where $D_1 = D_2 = \{c, d\}$ and the enumerated states of D are $D = \{1, 2, 3, 4\}$. This approach can be extended to utilize a state space which holds more information. To analyze the tit-for-two tat strategy, the state space is D^2 where the 16 elements of D^2 are enumerated by extending the numbering system from D by defining $1' = (11) = ((c,c), (c,c))$, $2' = (12) = ((c,c), (c,d))$, $\ldots$ $4' = (14) = ((c,c), (d,d))$, $5' = (21) = ((c,d), (c,c))$, $\ldots$ and $16' = (44) = ((d,d), (d,d))$.

According to Player 2's tit-for-two tat strategy, transitions of the type $((d,c), (c,c)) \to ((c,c), (*,d))$ are feasible. Player 2 plays "d" which Player 1 played 2 rounds ago, as

$$B = \begin{bmatrix}
1&0&1&0&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&0&1&0&1&0&0&0&0&0&0&0&0&0\\
0&0&0&0&0&0&0&0&1&0&1&0&0&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&1&0&1&0\\
1&0&1&0&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&0&1&0&1&0&0&0&0&0&0&0&0&0\\
0&0&0&0&0&0&0&0&1&0&1&0&0&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&1&0&1&0\\
0&1&0&1&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&0&0&1&0&1&0&0&0&0&0&0&0&0\\
0&0&0&0&0&0&0&0&0&1&0&1&0&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&0&1&0&1\\
0&1&0&1&0&0&0&0&0&0&0&0&0&0&0&0\\
0&0&0&0&0&1&0&1&0&0&0&0&0&0&0&0\\
0&0&0&0&0&0&0&0&0&1&0&1&0&0&0&0\\
0&0&0&0&0&0&0&0&0&0&0&0&0&1&0&1
\end{bmatrix}$$

FIGURE 2. Player 2 tit-for-two tat strategy

recorded by the state $((d,c),(c,c))$. Player 1's action is unspecified by this strategy of Player 2, thus the "$*$". Incorporating memory in the states of the bounded recall strategy implies that each state can only transition to one of 4 possible states (the 4 states that begin with (c,c)), not 16, as may initially be expected. This will be important when the complexity of strategies are measured. For the tit-for-two tat strategy, in terms of the symbols of D^2, $9' = (31) = \{(d,c),(c,c)\}$ transitions to $2' = (12) = \{(c,c),(c,d)\}$ or $4' = (14) = \{(c,c),(d,d)\}$. The (c,c) is shifted to the left from $\{(d,c),(c,c)\}$, while the $(*,d)$ comes from the action that Player 2 takes, allowing for each possibility for Player 1's action. Because states $2'$ and $4'$ are possible (under the tit-for-two tat strategy) when the history is given by state $9'$, there is a 1 in the 2^{nd} and 4^{th} columns of the 9^{th} row of the transition matrix associated with the tit-for-two tat strategy. The rest of the transitions for Player 2 under the tit-for-two tat strategy are given by the 16×16 transition matrix B, displayed in Figure 3.1. A complete list of all the states for, and their transitions under, the tit-for-two tat strategy appears in Appendix II.

The tit-for-two tat strategy defines a subshift $\Sigma \subset \Sigma_D$, but the subshift is not generated by any 4×4 transition matrix on the 4 one-step symbols of D. As described above symbol sequences beginning 3.12... and 3.14... are points included in the subshift Σ defining tit-for-two tat, but 3.11... is excluded from Σ by the rules of the strategy. The transition matrix B generates the subshift $\Sigma_B \subset \Sigma_{D^2}$ using the lift $l : \Sigma_D \to \Sigma_{D^2}$. For example: $l(3.12...) = (31).(12)...$ and $l(3.14...) = (31).(14)....$ It is more convenient in general to simply rewrite 2.13... as 21.3... emphasizing that two symbols are recorded at a time, and agree to interpret the overlapping symbols, each, as a new symbol in D^2.

Notice that Σ_{D^k} formally includes *all* possible sequences of the m^k symbols in D^k. The shift map on Σ_{D^k} cannot correspond to play on m^k states of D^k due to the use of labels in D^k to include the history of the past play. History cannot be chosen independently at each play, since the first $k-1$ terms are determined from the previous state. Therefore, transitions between each of the m^k labels in D^k can have at most only m outcomes each. This implies an important restriction to the size of a lifted full-shift, $l(\Sigma_D) \subset \Sigma_{D^k}$.

Proposition 3.1. *The full-shift Σ_D lifts to a subshift of Σ_{D^k}.*

Proof. Applying the shift map to points $\sigma \in \Sigma_D$ causes the k^{th} previous bit to be forgotten, while only one new bit from the set D is generated. The other $k-1$ bits are remembered, and therefore each of the m^k nodes in D^k can transition to only one of m possible new nodes. $\square$

Define a "randomized strategy" to be a bounded recall strategy of k rounds where every history maps to the probability distribution which gives equal probability to every action in D. The randomized strategy on the m states of D (not D^k) is defined by the $m \times m$ matrix R_m where every entry in the matrix is a 1. Note that the full-shift Σ_D is generated by a randomized strategy, $\Sigma_{R_m} = \Sigma_D$. Realize that the full-shift can be generated by any strategy that supports every possible action.

Example 3.2. Lifting the full-shift on 4 states.

Let $N = \{1, 2\}$ and $D = D_1 \times D_2$ as in Example 2.1. Hence, a randomized strategy is represented by the 4×4 transition matrix R_4 of all 1's. It follows that the full-shift $\Sigma_{R_4} = \Sigma_D$. The matrix R_4 lifts to the 16×16 matrix labeled R_4' in Figure 3, which, in agreement with the proof of Proposition 3.1, has only 4 nonzero elements per row. In general, define R_m' to be the lift of R_m to D^2. Consequently, $R_m^{(k)}$ is the lift of R_m to D^k. Proposition 3.1 dictates that $R_m^{(k)}$ has only m nonzero elements per row for all k. In Section 5, this property preserves the complexity of a given strategy under a lift.

Theorem 3.2. *Given a bounded recall strategy of k rounds on D, described by a subshift of Σ_D, the lift $l : \Sigma_D \to \Sigma_{D^k}$ yields an $m^k \times m^k$ transition matrix B which uniquely defines the strategy and generates the proper subshift $\Sigma_B \subset \Sigma_{D^k}$.*

Proof. All m^k possible permutations of k-step memories are represented by labels from D^k, and all possible transitions between D^k symbols can be represented by a $m^k \times m^k$ transition matrix. Such a matrix generates the subshift of all paths through the directed graph, which corresponds to all plays under the strategy. $\square$

$$R_4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, \quad R_4' = \begin{bmatrix}
1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\
1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1
\end{bmatrix}$$

FIGURE 3. "Randomized strategy for $N = \{1, 2\}$ and $D = D_1 \times D_2$ where $D_1 = D_2 = \{c, d\}$ and the enumerated states of D are $D = \{1, 2, 3, 4\}$ The lift of R_4 to $k = 2$ preserves the fact that at each play, the randomized strategy only has four possible outcomes: R_4' only has 4 nonzero elements per row.

While it is important to lift a k-step strategy, represented by a subshift of the symbol space, $\Sigma \subset \Sigma_D$, at least up to Σ_{D^k} to represent the strategy by a transition matrix B which generates the subshift $\Sigma_B \subset \Sigma_{D^k}$, there is no danger in lifting the strategy even higher. There is a unique $m^i \times m^i$ transition matrix A representation of the strategy with outcomes recorded in the subshift $\Sigma_A \subset \Sigma_{D^i}$ if $i \geq k$. Over-lifting the subshift is analogous to representing a 2-D plane as a plane in a higher dimensional space. The plane is still two dimensional, even though it is embedded in a higher dimensional space. In the next section, overlifting a given player's strategy is useful to compare strategies of players basing their strategies on different amounts of information, *i.e.*, players use bounded recall strategies that depend on a different number of rounds.

4. AN ALGEBRA ON TRANSITION MATRICES

Ultimately, we want to be able to measure the complexity of a single player's strategy, as well as the strategy profile accounting for the strategies used by all of the players. To proceed, we develop a technique to find the transition matrix associated with a strategy profile, if the transition matrices of the individual players are known. To utilize this technique, all players' transition matrices must be of the same size.

According to the previous section, it is possible to lift strategies and represent all the individual strategies by transition matrices of the same size. The strategies can be represented by $m^k \times m^k$ matrices where the bounded recall strategies are of round k or less for all players. This presupposes that the different players use the same symbol set.

Definition 4.1. For two transition matrices, A and B of equal dimension, define $C = A \wedge B$ as the matrix where the entry $c_{s,t}$ of C is equal to $a_{s,t} \times b_{s,t}$. The operation $\wedge$ is called the "and" operation.

Notice that the $(s,t)^{th}$ entry of $A \wedge B$ is 1 if both A and B have a 1 in their $(s,t)^{th}$ entries; otherwise, the $(s,t)^{th}$ entry of $A \wedge B$ is 0. To think of this operation in terms of what states are possible from the players' strategies, a transition to a new state is only possible if all players' strategies allow for that possibility. The new transition matrix C defines a subshift Σ_C which may further restrict what outcomes are possible between the two players. Accordingly, $\Sigma_C = \Sigma_A \cap \Sigma_B$. If the repeated game has n players, the n players' transition matrices are lifted to the same dimension (if necessary) and the strategy profile's transition matrix is the outcome of the $\wedge$ operation on the n lifted, transition matrices. The following example determines the transition matrix associated with the repeated Prisoner's Dilemma where Player 1 plays tit-for-tat and Player 2 plays tit-for-two tat.

Example 4.1. Tit-for-tat *vs.* tit-for-two tat in the Repeated Prisoner's Dilemma

From Example 2.1, Player 1's tit-for-tat strategy can be lifted and represented by a 16×16 transition matrix. This 16×16 transition matrix, utilizing the symbol notation from Example 3.1, is given in Figure 4. Player 2's tit-for-two tat strategy yields the matrix B in Figure 2, as determined in Example 3.1. The resulting transition matrix for the strategy profile of the two strategies is given in Figure 4.

To examine the complexity of a single player's strategy, simply take the $\wedge$ product of that player's strategy with the transition matrices for the other players' randomized strategies. Recall that the randomized strategy places equal probability on every action occuring; the matrix associated with the randomized strategy is given by R. As you would expect, this should not restrict what outcomes occur. This is represented by the following proposition which follows immediately.

Proposition 4.1. *Let the transition matrix A be associated with Player i's bounded recall strategy of k rounds. Then $A \wedge R_k = A$ and $\Sigma_{A \wedge R_k} = \Sigma_A$.*

5. Complexity and Topological Entropy

Defining complexity of both individual strategies and strategy profiles is quite natural using the symbolic dynamical description of players' strategies and the algebra developed in the previous section. The key is to compare the "size" of the subshift

$$A' = \begin{bmatrix}
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1
\end{bmatrix}$$

FIGURE 4. Player 1's Tit-For-Tat Strategy Lifted to D^2

$$C = \begin{bmatrix}
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0
\end{bmatrix}$$

FIGURE 5. The Strategy Profile Where Player 1 Plays Tit-For-Tat and Player 2 Plays Tit-For-Two Tat: $C = A' \wedge B$

which defines a given strategy by inclusion of all possible subsequences of play. By identifying the complexity of a strategy profile in terms of the associated subshift, we classify strategies in terms of their outcomes. This is in contrast to a classification scheme which may classify complexity in terms of rules or restrictions on possible outcomes. In fact, as the number of rules increases, the associated subshift, or outcome space, tends to get smaller.

Generically, other than a few trivial strategies with finite subshifts, the subshifts will tend to be uncountably infinite. Cardinality is not a sufficient measure to distinguish between two obviously varied strategies. As an example, notice that randomized play and tit-for-tat single player strategies both have uncountable subshifts, but should not be considered equally complex.

We measure complexity of a strategy or a strategy profile as the topological entropy of the associated subshift. Originally introduced by Adler, Konheim and McAndrew (1965) in the context of information theory, topological entropy has become a familiar tool in the theory of dynamical systems as a measure of the complexity of chaos. (The other thermodynamic function "metric entropy" also serves as an interesting measure of complexity in dynamical systems, but is not as readily calculated for the symbol dynamical description of strategies as developed in the previous sections.) Calculating topological entropy is particularly straight forward for dynamics of a subshift of finite type (when the transition matrix is of finite size).

To define the topological entropy of a subshift $\Sigma_A \subset \Sigma_{D^k}$, where A is the $m^k \times m^k$ transition matrix on the m^k symbols D^k of Σ_{D^k}, some additional definitions and concepts must be given. Define a word of length n as a particlar combination of n symbols from D^k: $\underbrace{(x_0, x_1, x_2, \ldots, x_{n-1})}_{n \text{ bits}}$ where $x_j \in D^k$. Thus, a point $\sigma \in \Sigma_{D^k}$ can be thought of as a word of infinite length, with a decimal added as the place holder, indicating the current round of play. The topological entropy h of a subshift Σ_A is the logarithm of the asymptotic growth rate of the words of length n found in the subshift, as n goes to infinity.

Definition 5.1. The word count of a subshift Σ_A is the number of subsequences of length n which are contained in the subshift, see Robinson (1995), and is denoted as

$$w_n(\Sigma_A) = \#\{(x_0, \ldots, x_{n-1}) : x_i = \sigma_i \text{ for some } \sigma = \sigma_0.\sigma_1\sigma_2 \cdots \in \Sigma_A\}.$$

Definition 5.2. The topological entropy of the subshift Σ_A is the scalar quantity $h(\Sigma_A)$, where

$$h(\Sigma_A) = \lim_{n \to \infty} \frac{\ln(w_n(\Sigma_A))}{n}. \tag{5.1}$$

Entropy is defined in terms of the logarithm of the growth rate. For most strategies, as n increases, there is an exponential explosion of the number of possible words. Entropy measures the exponent.

Theorem 5.1. *The range of the entropy function for a subshift $\Sigma_A \subseteq \Sigma_D$ is given by the inequality*

$$0 \leq h(\Sigma_A) \leq h(\Sigma_D) = \ln m.$$

Proof. For a subshift Σ_A, $h(\Sigma_A)$ is bounded below by 0 since $w_n(\Sigma_A)$ is positive. Since $\Sigma_A \subseteq \Sigma_D$, it follows that $w_n(\Sigma_A) \leq w_n(\Sigma_D)$ and $h(\Sigma_A) \leq h(\Sigma_D)$.

For Σ_D, the possible words of length n are all permutations of m elements n at a time, which equals m^n. Equation 5.1 implies that

$$h(\Sigma_D) = \lim_{n\to\infty} \frac{\ln(w_n(\Sigma_D))}{n} = \lim_{n\to\infty} \frac{\ln m^n}{n} = \ln m.$$

□

The full-shift on the m^k symbols of D^k is $h(\Sigma_{D^k}) = k \ln m$. If it were true that the full-shift Σ_D lifts to the full-shift Σ_{D^k}, then the entropy would not be a well defined measure of complexity. However, as stated in Proposition 3.1, the fullshift Σ_D lifts to a subshift, as required to preserve the entropy between representations of the same set of strategy outcomes.

Theorem 5.2. *Given a transition matrix A on the m symbols of D generating the subshift Σ_A, the lift $l : \Sigma_D \to \Sigma_{D^k}$ to a k-step recall representation with the $m^k \times m^k$ transition matrix B, preserves the entropy: $h(\Sigma_A) = h(\Sigma_B)$.*

The next theorem gives a useful technique to calculate the entropy of an arbitrary k-step bounded recall strategy. The proof appears in Robinson (1995).

Theorem 5.3. *Given a subshift of finite type Σ_A generated by the transition matrix A,*

$$h(\Sigma_A) = \ln \rho(A). \tag{5.2}$$

Example 5.1. From Figure 3, we calculate the largest eigenvalue of R_4 is $\rho(R_4) = 2$. Likewise, the spectral radius $\rho(R'_4) = 2$. As expected, the topological entropy is preserved by the lift and $h(\Sigma_{R_4}) = h(\Sigma_{R'_4}) = \ln 2$.

The restrictions that a strategy profile put on the sequence of play are at least as confining as the individual strategies in the profile. This is discussed below with regards to the $\wedge$ operation.

Theorem 5.4. *The topological entropy function is monotone nonincreasing under the "$\wedge$" operator. For two transition matrices A and B,*

$$h(\Sigma_{A \wedge B}) \leq h(\Sigma_A). \tag{5.3}$$

Proof. By definition of the "$\wedge$" operator $\Sigma_{A \wedge B} \subseteq \Sigma_A$. The rate of growth of words in the subset must be no greater than that in the larger set. Hence the result. □

Corollary 5.5. *A strategy profile has topological entropy less than or equal to the individual strategies that make up the profile.*

6. Comparing the Topological Entropy of Different Strategies

Comparing bounded recall strategies by examining their topological entropy gives a confusing account of complexity. It measures the complexity of the possible outcomes, not the complexity of the rules. In some sense, using each action with equal probability is not as complex as using a pure strategy which uses different actions depending on different histories. However, we use this measure to compare how varied the sequence of play can be given some initial string of information. First, we cite a lemma which aids in the proof of the theorems. The proof of the lemma appears in the Appendix.

Lemma 6.1. *An $s \times s$ matrix with t 1's in earch row and column, and 0's elsewhere has a largest eigenvalue of s.*

The first theorem establishes the intuition that a pure strategy profile is less complex than profiles that incorporate mixed strategies.

Theorem 6.2. *A strategy profile of pure strategies has topological entropy of zero.*

Proof. A strategy profile of pure strategies maps a state to a single state with certainty. Therefore, the associated transition matrix has a single 1 in every row. Consequently, the spectral radius of the transition matrix is 1. From Lemma 6.1 and Theorem 5.3, the topological entropy of a strategy profile of pure strategies is $\ln 1 = 0$. □

Individual players can use topological entropy to measure the complexity of their strategies, as well as the strategy profile. Pure strategies are again the simplest strategy, even for an individual who is uncertain about the strategy of the other players.

Theorem 6.3. *Suppose Player i uses a pure bounded recall strategy of k rounds. The topological entropy of this strategy is $\ln\left(\frac{m}{m_i}\right)$.*

Proof. The transition matrix associated with Player i's strategy is an $m^k \times m^k$ matrix where 1 appears in $\frac{m}{m_i}$ entries per row. Since Player i's strategy is pure, it restricts the transition to states that contain the action that Player i will take. Since Player i is uncertain about the other players' strategies, any of their actions are possible. There are $\frac{m}{m_i}$ such actions. It follows that the spectral radius of an $m^k \times m^k$ matrix with $\frac{m}{m_i}$ 1's in each row is $\ln\left(\frac{m}{m_i}\right)$ □

Realize that this measure of complexity only represents the complexity of the set of possible outcomes under a strategy or strategy profile. The following example gives three strategies whose possible outcomes have the same complexity, derived by their topological entropy. However, these three strategies are significantly different. Other measures of the complexity distinguish among these strategies.

Example 6.1. Comparison of Measures of Complexity

Let S_1, S_2, S_3 be three possible strategies for player 1 in the repeated Prisoner's Dilemma. Strategies S_1 and S_2 will be bounded recall strategies which depend only on the previous round. Let S_1 be the totally random strategy which selects C or D with equal probability regardless of the past actions. As seen previously, the spectral radius of the associated transition matrix is $\rho_{S_1} = \ln 4$. This value is the largest that the complexity can be for a 2 player game where both players have 2 strategies.

Let S_2 be defined by

$$\begin{aligned}
S_2((C, C)) &= p_1 C + (1 - p_1) D \\
S_2((C, D)) &= p_2 C + (1 - p_2) D \\
S_2((D, C)) &= p_3 C + (1 - p_3) D \\
S_2((D, D)) &= p_4 C + (1 - p_4) D
\end{aligned}$$

where $0 < p_k < 1$ for all k. This strategy is similar to S_1 since every action is possible after any sequence of play. And, $S_1 = S_2$ when $p_k = \frac{1}{2}$ for all k. Since every possible action can occur after any sequence of play, the transition matrix for S_2 is the same as the transition matrix for S_1. Therefore, S_1 and S_2 have the same complexity.

Let S_3 be a bounded recall strategy of 2 rounds. Define S_3 by

$$\begin{aligned}
S_3((C, C), (C, C)) &= q_1 C + (1 - q_1) D \\
S_3((C, C), (C, D)) &= q_2 C + (1 - q_2) D \\
S_3((C, C), (D, C)) &= q_3 C + (1 - q_3) D \\
S_3((C, C), (D, D)) &= q_4 C + (1 - q_4) D \\
S_3((C, D), (C, C)) &= q_5 C + (1 - q_5) D \\
S_3((C, D), (C, D)) &= q_6 C + (1 - q_6) D \\
S_3((C, D), (D, C)) &= q_7 C + (1 - q_7) D \\
S_3((C, D), (D, D)) &= q_8 C + (1 - q_8) D \\
S_3((D, C), (C, C)) &= q_9 C + (1 - q_9) D \\
S_3((D, C), (C, D)) &= q_{10} C + (1 - q_{10}) D \\
S_3((D, C), (D, C)) &= q_{11} C + (1 - q_{11}) D \\
S_3((D, C), (D, D)) &= q_{12} C + (1 - q_{12}) D \\
S_3((D, D), (C, C)) &= q_{13} C + (1 - q_{13}) D \\
S_3((D, D), (C, D)) &= q_{14} C + (1 - q_{14}) D \\
S_3((D, D), (D, C)) &= q_{15} C + (1 - q_{15}) D \\
S_3((D, D), (D, D)) &= q_{16} C + (1 - q_{16}) D
\end{aligned}$$

where $0 < q_k < 1$ for all k. Then, once again, every sequence of play is possible. The transition matrix for S_3 is a 16×16 matrix. The lifted transition matrix for S_1 and S_2 is the same as the transition matrix for S_3. Hence, $\rho_{S_3} = \ln 4$, also.

This measure of complexity does not distinguish among the three strategies S_1, S_2, and S_3. However, if $p_j \neq p_k$ for $j \neq k$ and $q_j \neq q_k$ for $j \neq k$, then the minimal number of states of finite automata representing these three strategies all differ. (An automaton with 1 state can mimic S_1, while automata with 4 states and 16 states, respectively, can implement S_2 and S_3.) The difference is one of perspective. Topological entropy measures the complexity of the possible outcomes, as opposed to the complexity of strategy implementation.

7. Conclusion

We have laid the ground work for the comparison of outcomes of strategies. Just how complex can the sequences of play be? Topological entropy is a natural way to measure this complexity. Some of the results are very satisfying, *e.g.*, the topological entropy of the transition matrix associated with a pure strategy profile is zero. However, other consequences are a bit disheartening, *e.g.*, different strategies have may have the same complexity of outcomes, but be implemented by automata with different numbers of minimal states.

The perspective of dynamical systems which considers the action space as an alphabet and the set of all possible sequences of play as a full shift on a symbol space is intriguing. As Neyman (1985) and Rubinstein (1986) examine how the complexity of automata affect what equilibrium outcomes are possible in repeated games, a similar course of research can be taken with this new measure of complexity. Should players prefer to minimize the complexity of the possible outcomes or maximize it? A player may want to use a strategy that has the same set of possible outcomes as other strategies. This could be beneficial since opponents would find it more difficult to determine his strategy.

Topological entropy only considers what can happen and does not weigh this outcomes with any probabilities. Other definitions of entropy, *e.g.*, metric entropy (which is similar in spirit to the strategic entropy defined by Neyman, 1996), may yield a satisfactory technique to consider these cases.

Appendix I: Details and Proofs

Proof. (of Lemma 6.1) One can check that the column vector v consisting of all 1's, is in fact an eigenvector. The equation $A(v) = \lambda(v)$ yields $\lambda = t$, since calculating each row is equivalent to the sum of t "1's", by hypothesis. However, it still needs to be shown that $\lambda = t$ is the largest eigenvalue.

The proof of the theorem is illustrated by the following example. Consider the given 4×4 matrix A, with two "1's" in each row and column, in the following eigenvalue equation.

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \cdot \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix} = \lambda \begin{bmatrix} a \\ b \\ c \\ d \end{bmatrix}. \tag{7.1}$$

Matrix multiplication yields two sets of equations:

$$\begin{aligned} a + b &= \lambda a \\ a + b &= \lambda b \Longrightarrow \lambda a = \lambda b \\ \\ c + d &= \lambda c \\ c + d &= \lambda d \Longrightarrow \lambda c = \lambda d \end{aligned} \tag{7.2}$$

The equation $\lambda a = \lambda b$ has two solutions, $\lambda = 0$ and $a = b$. Substituting $a = b$ into the first two equations yields $\lambda = 2$. Likewise, $\lambda c = \lambda d$ also implies $\lambda = 0$ and $c = d \Longleftrightarrow \lambda = 2$.

Consider the general $s \times s$ matrix A with t 1's per row and column. In fact, t must divide s evenly, which is easily visualized by permuting rows and columns so that the A is in $t \times t$ block diagonal form, with 1's along the diagonal blocks. If t does not divide s, then there will be overlap between the blocks, and some rows or columns will have more than t ones. Return A to its original arrangement. The eigenvalue equation

$$A(v) = \lambda(v) \tag{7.3}$$

consists of s/t blocks of equations analogous to those in equations 7.2. Re-indexing the v to v allows the order of $1 \leq j \leq N$ to be chosen so that each block of equations

$$v'_{cn} + v'_{cn+1} + \dots + v'_{cn+n-1} = \lambda v'_k \tag{7.4}$$

where $c \in 0, 1, \dots N/n$, and $k \in cn, cn+1, \dots cn+n-1$, is written with associated blocks in groups as in Equation 7.2. Re-indexing effectively block diagonalizes the problem. Hence, Equation 7.4 yields the s/t groups of equations,

$$\lambda v'_{cn} = \lambda v'_{cn+1} = \cdots = \lambda v'_{cn+n-1}. \tag{7.5}$$

Each group of equations has the trivial solution $\lambda = 0$ and $v'_{cn} = v'_{cn+1} = \dots = v'_{cn+n-1}$. Substitution of the nontrivial solution back into Equation 7.4 yields $\lambda = n$ for all c. $\square$

Appendix II: Transitions for Example 4.1

This appendix extends the list of transitions in Example 4.1 behind the derivation of transition matrix B, Figure 2. We explicitly list all the allowed transitions of the Player 2 bounded recall strategy of 2 rounds, tit-for-two-tat in the repeated prisoner's dilemna. The following are allowed transitions according to Player 2's strategy, "I do what you did two rounds ago."

$$\begin{array}{ccccccc}
((c,c),(c,c)) & \rightarrow & ((c,c),(*,c)) & \Longleftrightarrow & 1' & \rightarrow & 1' \text{ or } 3' \\
((c,c),(c,d)) & \rightarrow & ((c,d),(*,c)) & \Longleftrightarrow & 2' & \rightarrow & 5' \text{ or } 7' \\
((c,c),(d,c)) & \rightarrow & ((d,c),(*,c)) & \Longleftrightarrow & 3' & \rightarrow & 9' \text{ or } 11' \\
((c,c),(d,d)) & \rightarrow & ((d,d),(*,c)) & \Longleftrightarrow & 4' & \rightarrow & 13' \text{ or } 15' \\
((c,d),(c,c)) & \rightarrow & ((c,c),(*,c)) & \Longleftrightarrow & 5' & \rightarrow & 1' \text{ or } 3' \\
((c,d),(c,d)) & \rightarrow & ((c,d),(*,c)) & \Longleftrightarrow & 6' & \rightarrow & 5' \text{ or } 7' \\
((c,d),(d,c)) & \rightarrow & ((d,c),(*,c)) & \Longleftrightarrow & 7' & \rightarrow & 9' \text{ or } 11' \\
((c,d),(d,d)) & \rightarrow & ((d,d),(*,c)) & \Longleftrightarrow & 8' & \rightarrow & 13' \text{ or } 15' \\
((d,c),(c,c)) & \rightarrow & ((c,c),(*,d)) & \Longleftrightarrow & 9' & \rightarrow & 2' \text{ or } 4' \\
((d,c),(c,d)) & \rightarrow & ((c,d),(*,d)) & \Longleftrightarrow & 10' & \rightarrow & 6' \text{ or } 8' \\
((d,c),(d,c)) & \rightarrow & ((d,c),(*,d)) & \Longleftrightarrow & 11' & \rightarrow & 10' \text{ or } 12' \\
((d,c),(d,d)) & \rightarrow & ((d,d),(*,d)) & \Longleftrightarrow & 12' & \rightarrow & 14' \text{ or } 16' \\
((d,d),(c,c)) & \rightarrow & ((c,c),(*,d)) & \Longleftrightarrow & 13' & \rightarrow & 2' \text{ or } 4' \\
((d,d),(c,d)) & \rightarrow & ((c,d),(*,d)) & \Longleftrightarrow & 14' & \rightarrow & 6' \text{ or } 8' \\
((d,d),(d,c)) & \rightarrow & ((d,c),(*,d)) & \Longleftrightarrow & 15' & \rightarrow & 10' \text{ or } 12' \\
((d,d),(d,d)) & \rightarrow & ((d,d),(*,d)) & \Longleftrightarrow & 16' & \rightarrow & 14' \text{ or } 16'
\end{array}$$

For every possible play, according to the strategy, there is a corresponding "1" entry in the transition matrix, ($1' \rightarrow 1'$ or $3'$ implies the $B_{1,1} = B_{1,3} = 1$), and every other matrix entry is "0." That is, $B_{i,j} = 0$ implies that the transition $i' \rightarrow j'$ cannot occur in the given strategy.

References

Adler, R., Konheim, A., and M. McAndrew 1965. Topological Entropy. *Trans. of the Amer. Math. Society* 114.

Johnson, M.R. 1995. Algebraic Complexity of Strategy-Implementing Semiautomatons for Repeated-play Games. *mimeo.*

Kalai, E. 1990. Bounded Rationality and Strategic Complexity in Repeated Games. in *Game Theory and Applications*, T. Ichiishi, A. Neyman, and Y. Tauman ed., Academic Press.

Kalai, E. 1995. Games, Computers, and O.R. Center for Mathematical Studies in Economics and Management Science Discussion Paper #1141, October 1995.

Kalai, E. and A. Neme. 1992. The Strength of a Little Perfection. *Int. J. of Game Theory* 20: 335-355.

Kalai, E., Samet, D. and W. Stanford. 1988. A Note on Reactive Equilibria in the Discounted Prisoner's Dilemma and Associated Games. *Int. J. of Game Theory* 17: 177-186.

Kalai, E. and W. Stanford. 1988. Finite Rationality and Interpersonal Complexity in Repeated Games. *Econometrica* 56: 397-410.

LEHRER, E. 1988. Repeated Games with Stationary Bounded Recall Strategies. *J. of Economic Theory* 46: 130-144.

NEYMAN, A. 1985. Bounded Complexity Justifies Cooperation in the Finitely Repeated Prisoner's Dilemma. *Economic Letters* 19: 227-229.

NEYMAN, A. 1996. Strategic Entropy and Complexity in Repeated Games. *mimeo.*

ROBINSON, C. 1995. *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos.* Ann Arbor, MI: CRC Press.

RUBINSTEIN, A. 1986. Finite Automata Play the Repeated Prisoner's Dilemma. *J. of Economic Theory* 39: 83-96.

DEPARTMENT OF MATHEMATICAL SCIENCES
UNITED STATES MILITARY ACADEMY
WEST POINT, NY 10996
E-mail address: `bollt@euler.math.usma.edu` and `jones@euler.math.usma.edu`