

Restricted Non-cooperative Games

Seth J. Chandler

University of Houston Law Center
schandler@uh.edu

Abstract. Traditional non-cooperative game theory has been an extraordinarily powerful tool in modeling biological and economic behavior, as well as the effect of legal rules. And, although it contains plausible concepts of equilibrium behavior, it does not contain a theory of dynamics as to how equilibria are to be reached. This paper on Restricted Non-Cooperative Games inserts dynamic content into traditional game theory and thus permits modeling of more realistic settings by imposing topologies that restrict the strategies available to players. It uses *Mathematica* to show how the payoff array used in conventional game theory, coupled with these strategy topologies, can construct a "game network", which can be further visualized, analyzed, and "scored" for each of the players. The paper likewise uses *Mathematica* to analyze settings in which each player has the ability to engineer its own strategy topology and suggests other potential extensions of Restricted Non-Cooperative Games.¹

Keywords: non-cooperative game theory, Mathematica, law, Nash Equilibrium, game network, New Kind of Science, directed graphs.

1 Introduction

In conventional non-cooperative game theory, each player can see and can instantaneously select any element of its strategy set in response to the other players' strategy selections.[1] In real settings, however, the strategies available to a player at any given time will often be a function of the strategy it selected at a prior time.[2] It may, for example, be possible to change only one aspect of a strategy at a time. Alternatively, as in work earlier done by the author in "Foggy Game Theory,"[3] the strategies may be placed in some cyclic topology and only changes within some distance of the current strategy are permitted. Sometimes these constraints on the dynamics of strategy selection may be the result of external circumstances or cognitive limitations on the part of the player; other times they may be deliberately engineered by the player itself.[4] Either

¹ *Mathematica* code used to create the framework is available from the author on request.

way, however, the result is to overlay the strategies with a network connecting them (a topology) in which some strategies are connected and others are not.²

2 From Strategy Topologies and Payoff Arrays to Game Networks

The left panel of Figure 1, produced using *Mathematica*'s GraphPlot command, shows a sample strategy network s_D for an imaginary driver who has strategies labeled A through E (each strategy representing some combination, perhaps, of care and frequency). Notice that while the driver can continue the strategy of B, once it abandons that strategy it cannot return. This is the sort of realism permitted by this extension of conventional game theory. Notice further that the strategies differ in their ability to immediately access other strategies: D can access C and E immediately while A can access A and E immediately. Another player (perhaps a pedestrian) might, for example, have strategies labeled 1-3, again perhaps some combination of care when walking and frequency of walks. The pedestrian's strategy network s_P is shown in the right panel of Figure 1, in which the dashed arrows show the strategy connections.

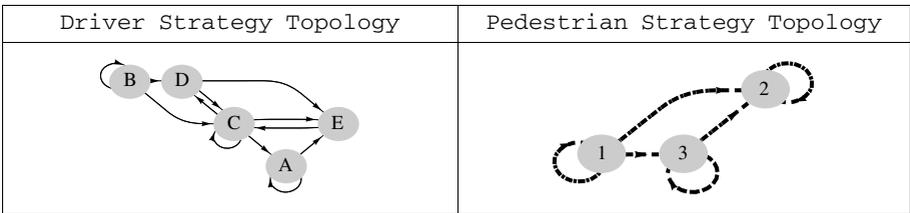


Fig. 1. Strategy topologies for driver and pedestrian

We can now use *Mathematica*'s structural operations to create a new network (directed graph or digraph) that is the Graph Cartesian Product of the networks s_D and s_P . Thus, given the strategy topologies shown in Figure 1, if the existing strategy combination is C1, the next strategy combinations could be A1, C1, D1 or E1 (if the driver moves) or C1, C2 or C3 (if the pedestrian moves).

In conventional game theory, the players get different payoffs depending on what strategy combination is selected. So too here. I assume that players are greedy (and perhaps not terribly clever) in that, in selecting their next move, they

² All strategies have at least one outgoing edge, though that edge can be a self-loop. Otherwise a player would not know what to do on their next move. One can imagine a yet more general case in which the strategies available to each player are a function not simply of the strategy employed by that particular player on the prior move but the strategy *combination* used on the prior move or move history. Conventional non-cooperative game theory may be thought of as a special case of restricted game theory in which each player has a complete graph as their strategy topology.

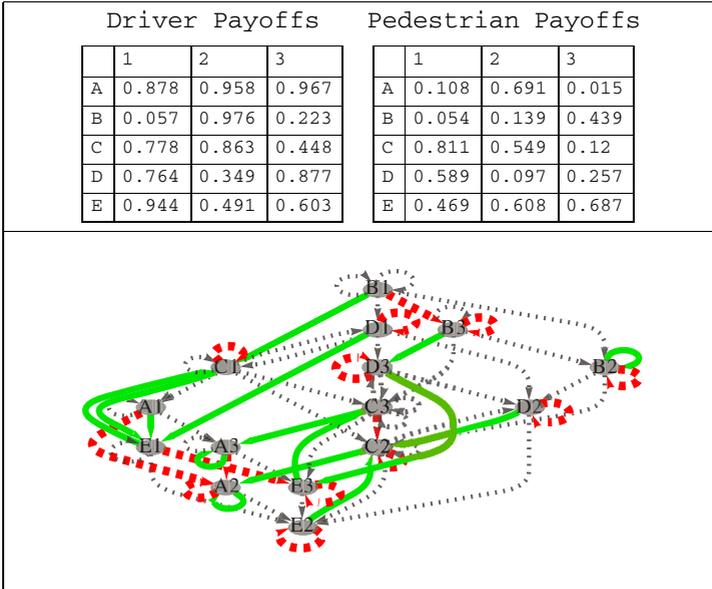


Fig. 2. Game network for driver and pedestrian based on payoff array

choose the one whose associated strategy combination (given the complementary strategy selections of the other players) offers them the highest payoff.³ This modification results in a thinning of the network created above so that, in an n -player game, only n edges can generally emerge from each node. Each edge represents the best move for one of the n players.

To use more formal notation, if there are n players in some restricted game and the strategy topology (graph) of player $i \in \{1, \dots, n\}$ is denoted as s_i , and the set of strategy combinations in the restricted game is $S \left(\bigotimes_{i=1}^n s_i \right)$, then the moves potentially available to player i from some strategy combination u may be written as Equation 1, where V is a function listing the vertices of a graph and E is a function listing the edges of a graph. One can then write the moves potentially available to all players from u as Equation 2 and the set of moves potentially available to all players from all strategy combinations as Equation 3.

³ More elaborate behavioral assumptions are certainly possible. Professor Steven Brams, for example, relies on a similar dynamic structure in his celebrated Theory of Moves.[2] He assumes, however, that the players can foresee and optimally negotiate an entire tree (acyclic directed network) of moves among strategy combinations. Restricted Non-Cooperative Game Theory avoids some of the issues associated with the construction of trees in the Theory of Moves in that, among other things, there are no "terminals." Instead, players confront a cyclic network of strategy combinations.

$$M[u, s_1, \dots, s_n, i] = \{ \{ \{ u \in S, v \in S \}, i \} \mid \{ u_i, v_i \} \in E[s_i] \wedge \forall_{j \neq i} u_j = v_j \} \quad (1)$$

$$M[u, s_1, \dots, s_n] = \bigcup_{i=1}^n M[u, s_1, \dots, s_n, i] \quad (2)$$

$$M[s_1, \dots, s_n] = \bigcup_{u \in S} M[u, s_1, \dots, s_n] \quad (3)$$

One now uses the payoffs to narrow the set of moves to a subset of the set of potential moves. A plausible way of doing so is, as in conventional game theory, to create a payoff function mapping each strategy combination to a set of payoffs, one for each of the players. One can conveniently represent that function as the $n \times |s_1| \times \dots \times |s_n|$ array P (dimensionality $n + 1$), where $|s|$ assumes its conventional meaning of the cardinality of a set s . P_i is in turn the n -dimensional array in which element $\{u_1, \dots, u_n\}$ represents the the payoff to the i^{th} player resulting from strategy combination u . One can then denote the restricted game G as having the following set of moves:

$$G[P, s_1, \dots, s_n] = \{ \{ \{ u, v \}, i \} \in M[s_1, \dots, s_n] \mid P_i[v] \geq \max_{\{ \{ u, m \}, i \} \in M[s_1, \dots, s_n, i]} P_i[m] \} \quad (4)$$

This mathematical formalism is visualized in Figure 2. The top panel represents the payoffs associated with each strategy combination in the sample driver-pedestrian game. The bottom panel shows the new "Game Network." Moves by the driver are shown as medium-width solid lines while moves by the pedestrian are thick dashed lines. Very faint, thin, dotted lines represent moves that were potentially available but were discarded because they were not the best move for either player.

3 Scoring Game Networks

We can now assign a score to each of the players from the game network described above. The player's score is simply a weighted average of the payoffs the player receives from each strategy combination. A plausible weighting scheme assumes that a random walk is taken by the players on the game network (starting from some random node) and that the weights are thus the stationary probabilities of being at each node (strategy combination). Thus, if there were, as in some conventional non-cooperative games, a single strategy combination to which all the players inexorably converged regardless of some "starting point" (a classic Nash Equilibrium), that strategy combination would receive a weight of one. If there were several such nodes, each would receive a weight corresponding to the size (number of nodes) of its basin of attraction. These weights can be computed readily by creating a Markov transition matrix associated with the game network and then computing the stationary values.⁴ The scores can be normalized by dividing each score by the score that would result if all nodes of

⁴ Mathematica has iterative constructs such as Nest or built-in functions such as Eigenvectors that make this process quite simple.

the network were weighted equally. In formal notation, the normalized score for player i in game $g = G[P, s_1, \dots, s_n]$ is equal to

$$\sigma[P, g, i] = \sum_{u \in V[g]} w_u[g] P_i[u] / |V[g]| , \tag{5}$$

where $w_u[g]$ is the weight accorded strategy combination u in game g . In our sample driver-pedestrian game, the normalized score is 1.6345 for the driver and 1.214 for the pedestrian.⁵

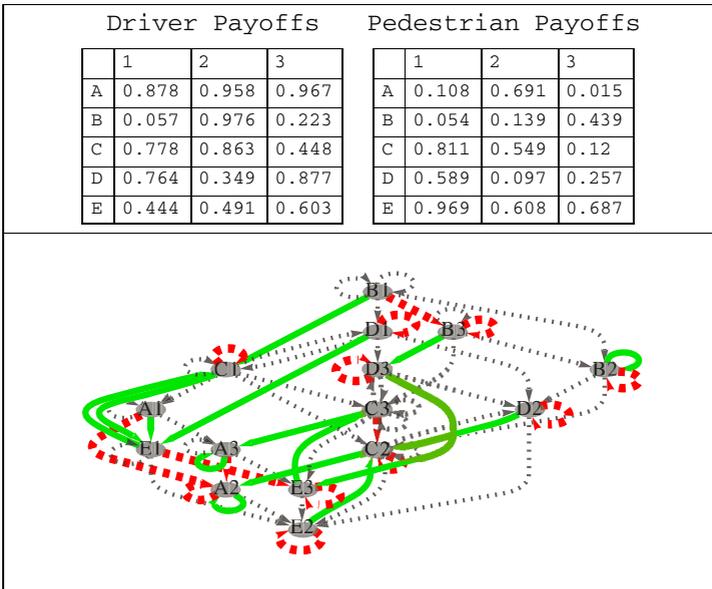


Fig. 3. Modification of payoff array generates new game network and scores

Just as in conventional game theory one can study how changes in the strategy combination-payoff mapping alter the Nash Equilibria and/or the payoffs at the equilibria, in restricted non-cooperative game theory one can study how changes in the strategy combination-payoff mapping alter the game network, which in turn alter the players' scores. Figure 3 shows the new payoff array (top panel) that results from requiring the driver to pay the pedestrian an amount of 0.5 if strategy combination E1 is employed, and the resulting new game network

⁵ In traditional game theory, the ability to find at least one "Nash equilibrium" for all games and the accompanying payoffs is preserved by permitting the players to use probabilistic strategies[1] and to then indulge contested assumptions about the behavior of players under uncertainty. Probabilistic strategy selection is not permitted here.

(bottom).⁶ The new normalized scores are 1.49 for the driver and 1.60 for the pedestrian.

4 A New Kind of Science Approach to Game Networks

With this framework in place we can also undertake studies not possible in conventional non-cooperative game theory. We can examine how, given a payoff matrix, changing the strategy topologies affects the associated game network, which in turn affects the weights received by each strategy combination, which in turn affects the scores received by each player. We can examine properties of the game network itself, such as the lengths of its maximal cycles. We can also, in effect, create a metagame in which the strategies are not just things such as A or B but also choices about whether to permit oneself to transition from A to B. Physical sabotage of otherwise existing transition possibilities can create such restrictions; so can economic relationships with third parties such that various strategy transitions become sufficiently unprofitable (dominated) and thus disregarded.

I now begin to examine this last proposition systematically using ideas of Stephen Wolfram's *A New Kind of Science* involving consideration of very simple cases and enumeration of possibilities.^[6] Consider a game of n players indexed over i with each player having $|s|$ strategies available to it. Each player now has $(-1 + 2^{|s|})^{|s|}$ possible strategy topologies. This is so because each strategy can connect with the other strategies in $2^{|s|}$ ways, but one of those ways, the empty set, is prohibited, as each player must always have a next move, and there are $|s|$ strategies to be considered. There are thus $\prod_{i=1}^n (-1 + 2^{|s|})^{|s|}$ strategy combination topologies that can exist. Although the number of strategy topologies can readily become quite large, if there are two players and each player has three strategies available to it, each player has 343 strategy topologies and there are 117649 strategy combination topologies that can be enumerated along with an identical number of associated game networks.⁷ It is well within the ability of the framework we have created and today's computers to extract the maximal cycle lengths and the scores for each of these game networks.

I can create a random payoff array and then create adjacency list representations for each possible strategy topology. I can then invoke a command that creates a losslessly compressed representation of the game network for all pairs of strategy topologies. On a 2006 Macbook Pro, the computation for all game networks of $n = 2, |s| = 3$ takes about 30 minutes and produces an expression consuming five megabytes. (*Mathematica* permits a monitor to be attached

⁶ Legal rules often change payoff arrays in just the fashion shown here by requiring one player in a "game" to pay another player in a "game" if certain strategy combinations are employed.^[5] Total payoffs are generally conserved within a strategy combination, unless there are transaction costs or payoffs to non-parties, in which event the total payoff can diminish.

⁷ If there were three players and each player had four strategies, there would be 129746337890625 game networks to be considered, which shows a limitation of a "brute force" approach on contemporary computers.

to the computation to watch its progress and the data can be stored for later examination.)

With this listing of all possible networks and a decompression algorithm, I can then compute the lengths of each maximal cycle for each of the 343^2 game networks. It ends up being most common for game networks to have the potential to cycle through at most three strategy combinations before returning to a previously played strategy combination. Maximal cycles of 7, 8 and even 9 strategy combinations prove possible, however. Indeed, we can focus on the smaller number of game networks that show complex dynamics with long cycles in which, depending on the sequence of player moves, there is at least the potential for many strategy combinations to be pursued before returning to a previously played combination.

Alternatively, I can take the game networks and compute a mapping between pairs of strategy topologies and the scores for each player. This computation takes only 10 minutes, much of which is spent in decompressing the game networks.

One can do several things at this point. One can examine which strategy topology tends to have the highest average score for a particular payoff array. Figure 4 shows the results of this experiment. It shows that both players tend to do best when they have significant choices available to them regardless of their current strategy choice. "Pre-commitment strategies," which attempt to restrict strategy selections, tend not to do well when one does not know the strategy topology of the opposing player.

Player 1					
Player 2					

Fig. 4. Strategy topologies yielding highest average payoffs

One can also examine the character of any pure traditional Nash equilibrium for the "meta-game" created by this process in which the "strategies" are now strategy topologies and the payoffs are the "scores." When one runs this experiment on the sample game shown above, one finds that there are eight Nash Equilibria.⁸ Figure 5 shows a sample Nash equilibrium.⁹

⁸ All the equilibria have the following characteristics: the second player always chooses strategy topologies in which it must move to strategy 1 no matter what strategy it has played before; the first player never permits a move to strategy 3 no matter what player 1 does and no matter what it has done on any prior move.

⁹ One could create meta - metagames by imagining that the players can not only alter their strategy topologies but also their ability to transition from among strategy topologies. Because this would create over 5.7×10^{70831} possible game networks, however, any exhaustive study of the possibilities is, for the foreseeable future, impractical.

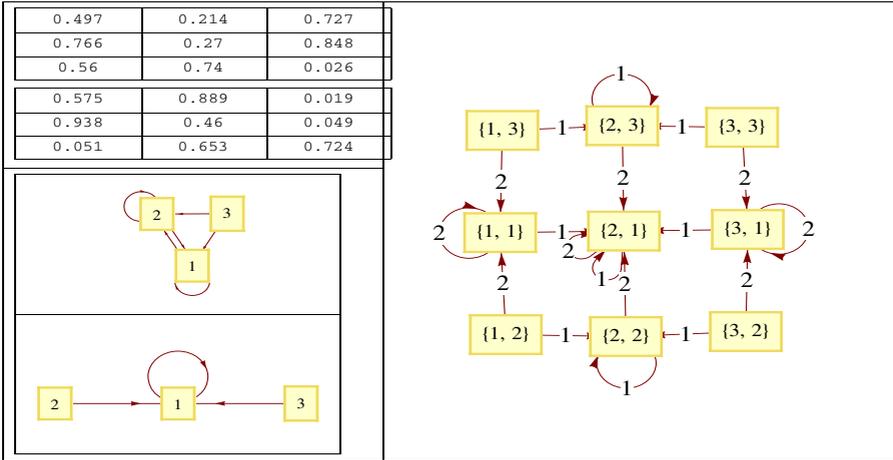


Fig. 5. A sample Nash Equilibrium set of strategy topologies and the associated game network

5 Conclusion

Mathematica successfully creates a useful and flexible framework for the study of n-player Restricted Non-Cooperative Games in which the players have potentially different strategy topologies. The paper does not purport to study this extension of game theory exhaustively. The intent is to develop a general set of tools from which further study can be profitably pursued.

Acknowledgments. The author thanks Jason Cawley of Wolfram Research, Professor Darren Bush of the University of Houston Law Center and Professor Steven Brams of New York University for extraordinarily helpful comments on a draft of this paper, as well as Yifan Hu of Wolfram Research for designing Mathematica’s GraphPlot package used extensively here.

References

1. Gintis, H.: Game theory evolving a problem-centered introduction to modeling strategic behavior. Princeton University Press, Princeton, N.J. (2000)
2. Brams, S.J.: Theory of moves. Cambridge University Press, Cambridge [England] ; New York, NY, USA (1994)
3. Chandler, S.J.: Foggy game theory and the structure of legal rules. In: Tazawa, Y. (ed.): Symbolic Computations: New Horizons. Tokyo Denki University Press, Tokyo (2001) 31-46
4. Dixit, A.K.: Strategic Behavior in Contests. American Economic Review **77** 5 (1987) 891-898
5. Baird, D.G., Gertner, R.H., Picker, R.C.: Game theory and the law. Harvard University Press, Cambridge, Mass. (1994)
6. Wolfram, S.: A new kind of science. Wolfram Media, Champaign, IL (2002)