

DRAFT COMPLEX BEHAVIOUR IN CHALLENGING SOCIAL SITUATIONS

DAVID GOFORTH AND DAVID ROBINSON

ABSTRACT. Social situations can be modeled as games with payoff structures. The complexity of the situation emerges from the pattern of payoffs. Inferring the outcome of the game from the payoff matrix is the conventional approach to describing, predicting or prescribing behaviour of participants. This approach makes a strong assumption of rationality of the players and suffers from other drawbacks as well. While Nash [8] proved a mixed equilibrium always exists, this appealing universality comes at a price. Mixed equilibria are not a natural and credible model of decision-making behaviour. Indeed, it has been shown that efficient algorithms to compute mixed equilibria likely do not exist for all games. [3] The pure Nash equilibrium, a more appealing model of behaviour, fails on universality: a subset of games - and a growing proportion as the number of players and choices increases - do not possess a pure Nash equilibrium.

Miller [5],[6] describes unpublished research in which he took a direct approach to investigating the complexity of player behaviour. In a series of evolutionary experiments on all the strictly ordinal 2×2 games, he quantified the complexity of behaviour of the players that succeeded.

Working from Miller's original data, we present a systematic inverse treatment of the complexity of social situations on the premiss that complex behaviour implies the game is challenging. Applying the topological map of the 2×2 games that we have developed[9], we demonstrate how complexity is distributed in payoff space and relate the complexity to readily identifiable features of the games.

1. COMPLEX SITUATIONS AND BEHAVIOUR IN 2×2 GAMES

Modeling a social situation as a game is done with the expectation of learning something about the features of the situation and how the participants react to them. Classical analysis takes the payoff structure of the game as given and attempts to infer the strategic choices of the players and the outcome of the interaction. The players are assumed to be rational, meaning they are presumed to have (i) the goal of maximizing their own

Date: May 18, 2009.

Key words and phrases. genetic algorithm, strategic complexity, game theory, ordinal games.

payoff and (ii) a strategy to achieve the goal. One modest rational goal is avoidance of the worst payoffs and there is a universally effective algorithm, Maximin or selecting the set of outcomes with the greatest minimum payoff, for achieving it. For any more ambitious interpretation of the rational goal, the prescription of a universal strategy for all games is problematic.

The problem is evident even in the simplest of models, the strict ordinal 2×2 games which we will use throughout this paper. The most widely known and applied solution concept, the pure Nash equilibrium, describes what a rational player's goal outcome might be but makes no assertions concerning how the player should achieve it. For some 2×2 games, there is no outcome defined and for others, there are two Nash equilibria. For a few games, including the infamous Prisoner's Dilemma, the Nash equilibrium identifies an outcome that is Pareto-dominated, clearly at odds with goal of rational players to optimize their payoffs. When the number of players or choices increases, the proportion of games where the Nash equilibrium fails also increases.

The mixed Nash equilibrium is a related approach that is universally applicable and specifies both the goal and behaviour of a rational player, however the behaviour is specified probabilistically and the optimum payoff is an expected value. The main drawback however is that the determination of the probabilities is demonstrably complex () and stretches the assumption of player rationality.

Part of the problem in defining solution concepts comes from the definition of the games themselves. While the quantification in terms of strategies and payoffs is universally applied, the procedure for play is less clear and more variable. What information do players have? Can they collaborate? Is play concurrent or sequential? Is the game played once or repeatedly? When the rules of play are explicitly specified, it is possible to propose better solution concepts, taking advantage of the additional constraints.

Nash proposed a solution concept, the Nash bargaining solution[7]) whereby players negotiate to an agreed outcome. This approach which assures the outcome is in the Pareto set and maximizes the combined payoff depends on collaboration between the players and assumes repeated play to achieve long term mixed outcomes. The Kalai-Smorodinski solution concept[4], under the same conditions, produces a result that assures both players achieve the same proportion of their maximum possible payoff.

Axelrod's famous Prisoner's Dilemma tournaments () were conducted under the more challenging conditions of concurrent play without collaboration. He did retain repeated game play. The most significant departure from the solution concepts already described however was the rethinking of the player. The rational goal of maximizing individual payoff was redefined and the assumption of an optimal rational tactic for achieving it was

dropped. Instead each player was an individual with a distinct algorithm for playing the PD game and a goal of maximizing total payoff against all players in the tournament. In addition to the payoff information, a player has access to the results of previous games and is permitted to include that information in determining the next strategic choice. Clearly, any results achieved under such rules are context dependent - a player's success is determined in part by the other players in the tournament.

In 1990, Miller extended the tournament model in two ways. (Miller 1991, 2006) First, instead of including a population of fixed players, he generated a set of random players based on an abstract model and allowed them to learn, via genetic algorithm, to play as they participated in repeated tournaments. With the populations that emerged, he took first steps toward identifying characteristics of successful behaviour. He operationalized the notion of complexity of playing strategy based on his abstract model and computed the relative complexity of strategies in the populations.

Miller's second extension to the tournament framework was to apply the approach to the entire class of strict ordinal 2×2 games. Based on the complexity of the most successful row and column players of all the games, he made some observations on which games are hard to play. To the extent that these games model social situations, we can claim further that the social situations also are challenging.

It is for exactly this kind of analysis that we have proposed an alternative organization of the ordinal 2×2 games based on a formal topological model. (Robinson and Goforth, 2005) This model provides a transformation operator that explicitly defines how similar games are to each other. A layout of the games according to this similarity measure, which we have named a periodic table, allows us to display data like Miller's complexity measures, identify clusters and hence determine connections to common game properties.

In the next section, we introduce the model of the 2×2 games and compare it to other organizations (Rapoport and Guyer 1966, 1976, Brams 1992). We also demonstrate the organizational power of the model by discussing the distribution of some common game sets such as coordination games, dominance solvable games, cyclic games, zero-sum games.

Next we briefly summarize Miller's results and observations before displaying them on the periodic table...***** (more here later)

2. VISUALIZING 2×2 GAMES ON THE PERIODIC TABLE

The periodic table (RandG 2005) was designed as an alternative organization of the 2×2 ordinal games with the goal of locating similar games close to each other on the premiss that global patterns might emerge if data

	L	R		L	R		L	R		L	R		L	R		L	R
U	1	3	2	3	2	3	1	2	2	1	3	1	2	3	3	4	
D	2	4	1	4	1	4	2	4	3	4	3	4	3	4	3	4	
	1		2		3		4		5		6		6				

FIGURE 1. Patterns of Row's Payoffs

gathered across many games were displayed on the table. In the periodic table, similarity is determined by comparing payoff matrices and formalized by defining operators that transform one game into another. We hypothesized that properties are related to underlying features of payoff structure and that clusters of related games would exhibit similar properties. This has turned out to be the case.

2.1. The structural logic of the periodic table. Under the usual assumption of equivalence under the exchange of choices, a player of a strict ordinal 2×2 game sees one of only six possible patterns of her own payoffs. These are enumerated for the Row player in Figure 1 and indexed as we will use them later. Likewise she faces one of six possible configurations of her opponent's payoffs. For a specific pair of payoff patterns, there are four ways to combine them to create distinct games. In all there are $6 \times 6 \times 4 = 144$ strict ordinal games. This set of games is portrayed on the periodic table.

Similarity of games is defined structurally by comparing payoff matrices. For any game, those most similar to it are the games that differ only in that a pair of ordinally consecutive payoffs (1 and 2, 2 and 3, or 3 and 4) for one player are in exchanged positions in the matrix. Hence, every game has six adjacent neighbours and we can represent the ordinal 2×2 games as a graph of 144 nodes with each connected by 6 edges to other game nodes¹.

This graph of 144 nodes and 432 edges cannot be mapped onto a plane so displaying it requires some decisions about which adjacencies to highlight. We do so based on the presumption that swapping high-valued payoffs 3 and 4 will alter the character of a game more than swapping low-valued payoffs 1 and 2 or 2 and 3. By concentrating on the four low-valued swaps, the periodic table will show games near those to which they are generally most similar.

The resulting graph with nodes of degree four is now disconnected into four simple toruses that are easily displayed as 6×6 'layers' of games.

¹We do not reduce the set to 78 games by assuming equivalence under the exchange of player roles. The adjacency based on 'swapping' a pair of payoffs is therefore unconstrained and the graph naturally includes all games and edges. Topologically, the graph is a 37-hole torus.

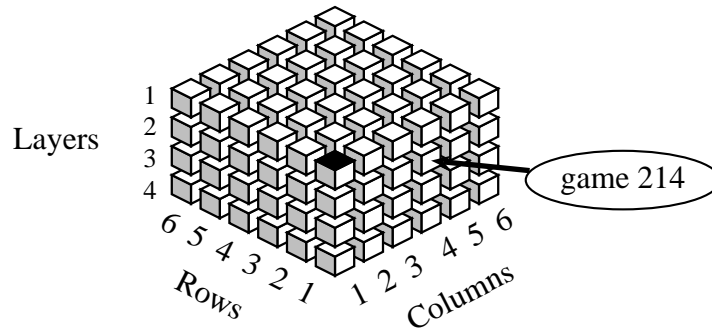


FIGURE 2. Arrangement of the 2×2 games by indices

Based on the initial observation that there are six patterns for each player, we can identify the rows of each layer with particular row player patterns and the columns with column patterns and thus index the games according to array position. If the four layers are piled up so that row patterns and column patterns coincide, then the four games with the same row and column index form a ‘stack’. In every stack are the four distinct games constructed from the same row and column payoff patterns. To each game is assigned a three digit index designating its layer, row and column. For example, game 214 in Figure 2 is on the second layer in row one at column four. Figure 3 shows the stack of four games, including game 214, that share the same row and column payoff patterns.

As well as payoff matrices, Figure 3 shows the alternative representation in payoff space, called an *order graph*, that is employed in the periodic table. Row payoffs are plotted on the horizontal axis with column payoffs vertical. Each point represents an outcome and the edges represent inducement correspondences (Greenberg**) corresponding to row’s choices (double) and column’s (single) on the payoff matrix.

In the periodic table, displayed in Figure 13 at the end of the paper, the four layers are arranged in a two by two array creating a 12×12 display of games. Clearly, the configuration is still underconstrained as there is choice in where to cut and project the toruses to form layers and there is choice in how to place the four layers. In the static printed version of the table, we have made specific decisions but in the dynamic online version, the user can




FIGURE 3. A stack: four games with the same payoff patterns for each player. The open dot indicates a Nash equilibrium outcome.

move, roll and reflect the layers to feature different adjacencies, including the swaps of 3 and 4 values that connect games on different layers².

2.2. Emergent properties of the periodic table. Although the adjacency of games on the periodic table is defined in strictly structural terms, there are many basic properties, including some used in defining taxonomies, that show up as clustered regions of the table.

2.2.1. Player patterns. The six patterns for one player form a closed cycle under consecutive swaps of 1 and 2 then 2 and 3. The three patterns that contain a dominant strategy are adjacent (indices 1, 5 and 6) in Figure 1 as are the three that do not have a dominant strategy (2, 3 and 4). The pattern labeled 1 is the weakest of the dominant strategy patterns, the only one with a dominant payoff of 2.

²For a complete treatment of the periodic table, see Robinson and Goforth(2005). A printable pdf file of the static table is available at (my url) as well as an online beta version of the interactive table.




FIGURE 4. Symmetric games on the diagonals of layers 1 and 3

2.2.2. Symmetry and Player Exchange. In a symmetric game, both players have the same payoff patterns, but only two of the four relative orientations produce a symmetric game. The twelve symmetric games are on layers 1 and 3, occupying the shaded positive diagonal in Figure 4 (replicating Figure 13 where the row and column indices are equal). Games with the same row and column index on layers 2 and 4 share the same pattern but are not symmetric.

The games that are equivalent under an exchange of players can be paired by reflecting across this diagonal on the full periodic table of Figure 4. For example, game 214 is equivalent to game 441 and 326 is equivalent to 362.




FIGURE 5. Layer Quadrants based on Dominance

2.2.3. *Combinations of patterns.* Each layer can be arrayed with the row player's dominant strategy patterns in the top three rows and the column player's dominant strategy patterns in the rightmost columns. All the figures in the paper are configured this way. Where these patterns meet in the top right corner are nine dominant strategy games, shown with the darkest shading in Figure 5. The nine in the top left are dominance solvable because the row player has a dominant strategy as are those in the bottom right because the column player has a dominant strategy pattern. In the bottom left (white in Figure 5) are the games with no dominant strategy. These descriptions apply to all layers.

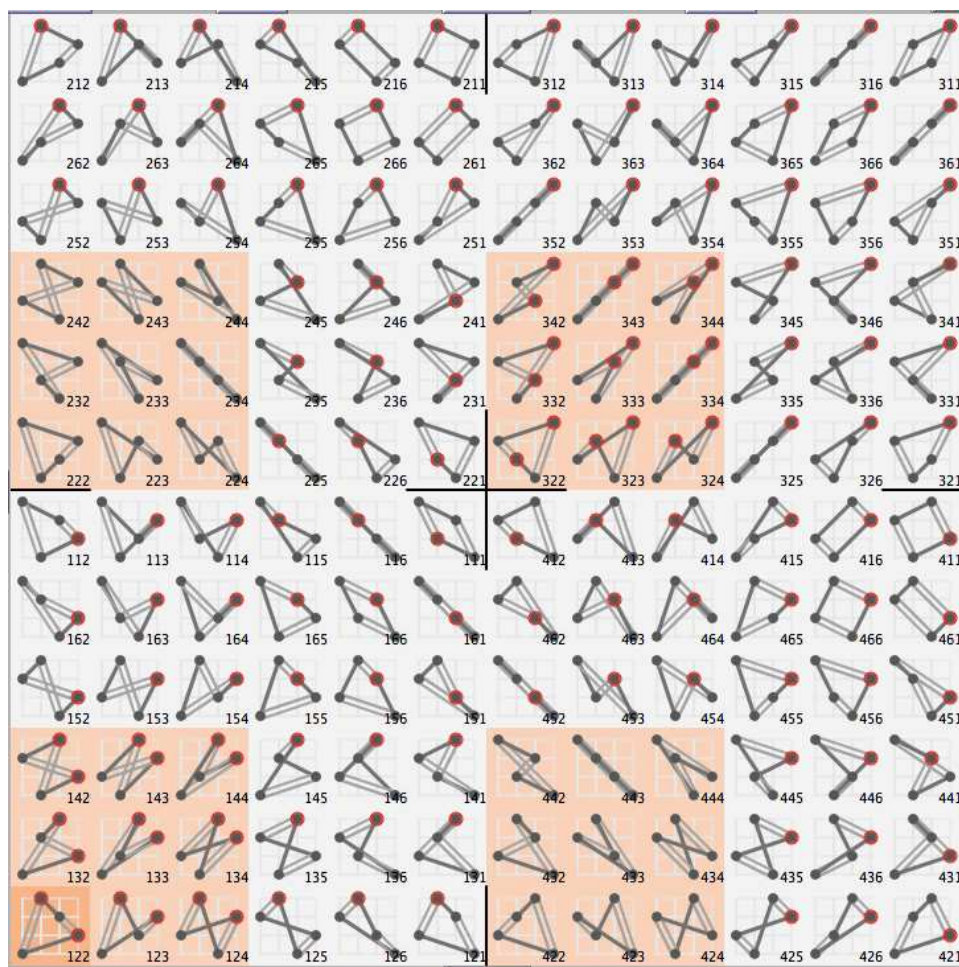


FIGURE 6. Quadrants with Zero or Two Nash Equilibria. Nash equilibria outcomes are circled in colour in all games.

2.2.4. *Equilibria.* All dominant strategy and dominance solvable games have a single Nash equilibrium. All the games with zero or two equilibria are composed of player patterns without a dominant strategy. These are the games in the bottom left quadrant of each layer, shaded in Figure 6, where both players have patterns 2, 3 or 4. This figure shows a circle around the Nash equilibrium outcomes. The games on layers 2 and 4 have no equilibria. Those on layers 1 and 3 have 2 equilibria but they are further distinguished. On layer three, all the games have one equilibrium that Pareto-dominates the other, the so-called coordination games including Stag Hunt (game 322); on layer one are the Battle-of-the-Sexes games, including Chicken (game 122), with two undominated equilibria.




FIGURE 7. Rapoport and Guyer Taxonomy on the periodic table

2.2.5. *Rapoport and Guyer's Taxonomy.* In Figure 7, the games are coloured and indexed according to the taxonomy in Rapoport and Guyer (1966). Reflected asymmetric games have the same index. Except for the blue shaded games (67, 68, 69) “two equilibria with non-equilibrium solution”, all of the taxonomic categories are contiguous in the periodic table. Some of the adjacencies are 3-4 swaps not shown in the configuration of Figure 7. Note that the “no equilibrium” category (purple) on layers 2 or 4, is exactly the set identified as having no dominant strategy for either player in Figure 5.

2.2.6. *Other patterns to demo?*

3. MILLER'S EXPERIMENTS

John Miller (1991) speculated that players of 2×2 games need more complex strategies to play some games successfully than others. To operationalize strategic complexity he defined playing strategies as Moore finite state machines with 16 states. He created a pair of populations of random machines and engaged them in iterated play of a 2×2 game in a tournament between the two populations. The results of the iterated play were used as a fitness measure for a genetic algorithm that evolved each of the populations as they 'learned' to play each other over a sequence of tournaments. The experiment was repeated for all 78 2×2 ordinal games.

As Miller postulated, in some games, the Moore machines evolved to include fewer states defining simple strategies. In other games, the machines retained more states to define more nuanced strategies. Miller tabulated results for each game once the conditions stabilized.

- The behaviour of the strategies was indicated by the relative frequency of the four outcomes in iterated play. This was an indicator of the success of the strategies.
- The complexity of strategies for row and column players was summarized as a pair of numbers in a normalized distribution of the average number of states in the final populations. On this scale, the simplest strategy was coded as -1.7 and the most complex as 1.6.
- Miller observed that encounters in some iterated games settled into patterns where the relative frequency of outcomes was not fixed but cycled through a repeating sequence. To measure the stability of the outcome frequencies, he created a statistic, a measure of temporal deviation, based on the standard deviation of moving averages. He tracked the maximum standard deviation of the outcome frequencies and the aggregate of the four, again normalizing the distributions. Aggregate temporal deviation ranges from -0.63 for stable distributions to 5.45. Miller observes that there is some correlation ($\rho = 0.35$) between strategic complexity and temporal outcome variation.

In Table 3 at the end of the paper, a few records from Miller's data are displayed.

Miller did not attempt to make a causal link from the features of games to the characteristics of player strategies but he did observe some apparent connections:

- (1) The games with dominant strategies have simple player strategies, except for the Prisoner's Dilemma. Games without Nash equilibria have complex strategies.

- (2) The Rapoport and Guyer taxonomic categories, except for the “no equilibrium” category (purple in Figure 7, games indexed 70-78), do not appear to correspond to the patterns of strategic complexity.
- (3) A small change in payoffs (a symmetric 1-2 swap for both row and column) makes a difference in the frequency of the most efficient (4,4) outcome in the two symmetric coordination games. These games 333 and 344 are included in Table 3 with Rapoport and Guyer indices 60 and 63.

4. MILLER’S RESULTS ON THE PERIODIC TABLE

Since the patterns of dominance solvability and Nash equilibria emerge so clearly on the periodic table, it seems justifiable to look for trends in other data when displayed on the periodic table. One could argue that a first analysis based on such a display is warranted, whether trends emerge or not. Even a chaotic display would be revealing as it would eliminate correlation with the dominance and equilibria and remove them as explanatory factors in the phenomenon under study. We have used Miller’s complexity information as a case study.

To display Miller’s data on the periodic table, we must map from 78 games to 144. The 12 symmetric games transfer directly but each of the 66 asymmetric games must be mapped to both variants. For one game, the row player will be interpreted as the row player and for the other, row will become the column player. The advantage is that now the player complexity measures of all 144 situations that a player of 2×2 may face can be directly compared.

The transformation must be done with a caveat however. The normalized complexity measures can only be put together under the assumption that the original complexity values came from the same distribution. Moreover, the variation in results between the two reflected versions of an asymmetric game, that would occur if a 144×144 experiment had been conducted, is lost. The possible variation is signaled by the difference in complexity of the players in the symmetric games. For example, the row player of Prisoner’s Dilemma has complexity 0.9 while the column player is a simpler -0.1. The distinct values for the players of symmetric games have been included in the data set.

4.1. Individual payoff patterns. Figure 8 shows the complexity measures for the row player on the periodic table. The simple strategies are red (complexity ≤ 0) and the complex ones black. To focus attention on the extrema of the distribution, the radius of the circle is proportional to the absolute value of the normalized complexity. The horizontal stripes suggest a connection between the complexity and the payoff pattern the row player faces.




FIGURE 8. Row's strategic complexity: red is simple, black is complex. Row patterns indicated at left.

Dividing the games into six groups of 24 based on row player pattern allows us to conduct pair-wise t-tests among the groups. The results in Table 1 confirm that the three groups associated with patterns without a dominant strategy (2, 3 and 4) are indistinguishable and are associated with complex strategies. Two of the dominant strategy patterns (5 and 6) are also indistinguishable and the strategies they evoke are simple. The other pattern, 1, is the weak dominant strategy pattern with intermediate complexity.

Applying the pairwise t-test to the games grouped according to column pattern finds no distinction in row complexity among the six groups. When the games are grouped by layer, there is a only significant difference, at the 5% level, between layers 1 and 3.

4.2. Opponents' payoffs: Interaction of payoff patterns. Although the column pattern was not found to be a primary indicator of row strategy complexity, it did prove to be a secondary factor. The row complexity of each game was represented as a deviation from the average complexity of its row pattern group. The resulting values were grouped according to column

TABLE 1. Pairwise t-test of Normalized Row Complexity aggregated into six groups of 24 games according to Row payoff pattern.

	2	3	4	5	6
1	0.004	0.000	0.002	0.000	0.000
2		0.508	0.898	0.000	0.000
3			0.594	0.000	0.000
4				0.000	0.000
5					0.878

Average complexity by row pattern group:

1	0.079	4	0.671
2	0.646	5	-1.104
3	0.763	6	-1.129

TABLE 2. Pairwise t-test of Residual Normalized Row Complexity aggregated into six groups of 24 games according to Column payoff pattern.

	2	3	4	5	6
1	0.007	0.000	0.001	0.570	0.704
2		0.326	0.401	0.001	0.019
3			0.865	0.000	0.001
4				0.000	0.001
5					0.336

pattern and a pairwise t-test was applied again. This time the three dominant strategy groups were significantly different from the three without dominant strategy, all but one at the 1% level. See Table 2. Thus the games we might expect to be the most challenging to play have row patterns 2, 3 or 4 and column pattern 2, 3 or 4 also. Of these 36 games, 35 are above the mean in complexity for the row player³.

Miller’s observation about the complexity of games in the Rapoport and Guyer “no equilibrium” category can now be put in context. “No equilibrium”, shaded purple in Figure 7, consists of nine games where there is no Nash equilibrium, precisely the games on the bottom left quadrant of layer 2 or 4 where neither player has a dominant strategy as in Figure 5. These are among the games predicted to inspire complex row strategies by their row patterns and column patterns.

³A secondary analysis based on layers found the only significant difference in row player complexity, in pairwise t-tests, between layers 1 and 3, as before.

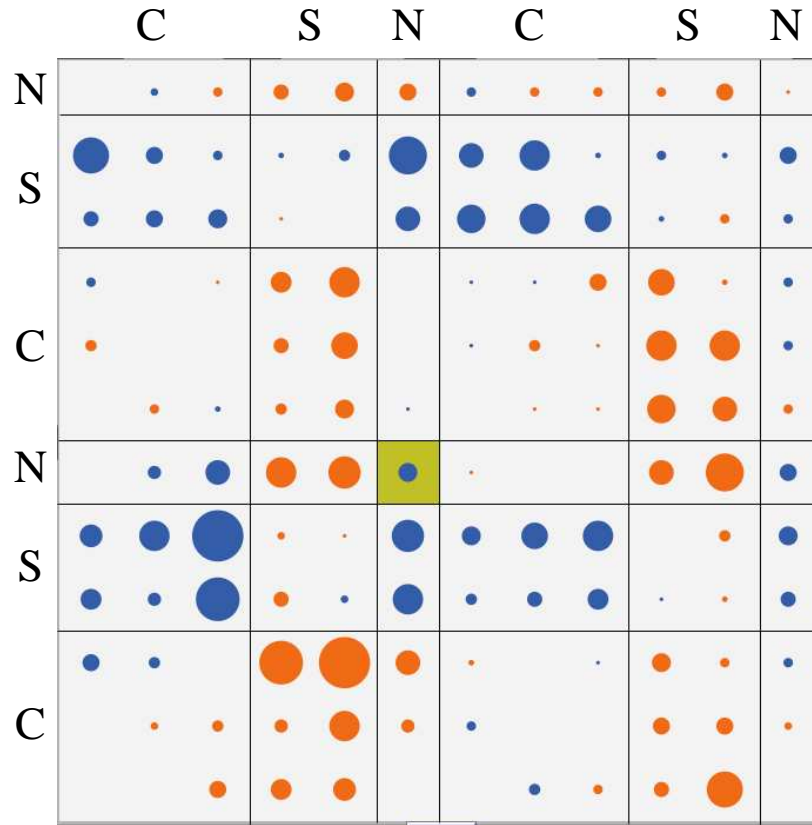


FIGURE 9. Differences in Complexity: Orange indicates the Row strategy is more complex

4.3. **Asymmetric games: when the challenges are not equal.** The analysis just described can be applied equally to the column strategies which, except for the symmetric games, are the same data set. By comparing the row and column complexities, we can make some estimate of equity or inequity of challenge faced by two players in a game. Considering only the primary effect of the payoff pattern each player faces, playing a game where the pattern is 5 or 6 is “simple” or “S”, playing a 2, 3 or 4 is “complex” or “C” and facing the pattern 1 is “neutral” or “N”. Each game can be predicted to be more or less equitable depending on the pattern each encounters.

How well does the data match these predictions? In Figure 9, the differences between the complexity measures of strategies opposing each other in every game are displayed. Where the row player has more complex behaviour, the dots are coloured orange; where the column player is more complex, the dots are blue. Dots ‘disappear’ when the complexity values are equal.

Along the diagonal of each layer, the squares labelled CC, SS, NN are populated with games where, predictably, the differences in complexity are small. The largest differences should be expected in regions labelled CS or SC and again, several games meet this expectation. Perhaps the most important use of this diagram however is to identify the aberrations. For example, the regions CN and NC adjacent to the Prisoner's Dilemma game 111 (green) are, unexpectedly, equitable in complexity. Figure 8 shows that Row's strategies in the pattern 1 row are relatively complex for those games (412, 413, 414). These are the three asymmetric games that possess Pareto-dominated dominance-solvable Nash equilibria and, as we will see later, a more complex row strategy has improved the payoff for the players⁴.

4.4. Looking for trends. Miller's observations concerning his complexity data came from examining the data and game characteristics he tabulated as in Table 3. Portraying the data on the periodic table reveals other patterns that can be examined more rigorously by other means. This has been demonstrated with the relation between complexity and payoff pattern in Section ???. Here we draw attention to some other trends in Miller's data but, in the interests of brevity, omit further analysis.

4.4.1. Dominance and Not. The weak dominant strategy pattern, 1, is adjacent by 12 and 34 swaps to the non-dominant pattern, 2, with only one peak payoff outcome. Call this the 'soft' boundary in the sense that the differences across it are not extreme and there might be some mutual bleeding of characteristics. The other dominant-nondominant boundary occurs between a strongly dominant pattern 5 and a nondominant pattern 4 with two opposite peaks of 4 and 3. We can hypothesize that there will be a more extreme disruption in game characteristics crossing this latter boundary than the former so call this a 'hard' boundary. Specifically we would expect to see more extreme differences in strategic complexity.

As an example, the most extreme changes in complexity occur across this boundary between games 114, and 115, 164 and 165, 154 and 155, a transition that changes dominance solvable games to games with two dominant strategies. The games 164 (Protector 4) and 165 exhibit the most extreme difference in strategic complexity for one player. Row's complexity measure is simple and almost fixed, -1.4 and -1.5. Column's strategy for 164 is 1.3 but for 165 it is reduced to -1.9. The relation between these two games is a C23 swap changing a dominant strategy pattern (5) to a pattern (4) without a dominant strategy. The swap also moves the Nash equilibrium from one outcome to another. In the intermediate boundary game (where payoffs 2

⁴The inverse argument applies for the games in the vertical region, games 221, 231, 241.

and 3 are equal) the Column inducement correspondence becomes a "don't care" choice. Hence, it is Column's pattern that changes but Row's payoff that is affected, 4 in 164 and 3 in 165.

Games 154 and 155 as well as 114 and 115 show similar changes in complexity for the Column strategy and fit the same structural patterns as 164 and 165. In 115, Row's payoff at the Nash equilibrium is only 2 and the outcome is inefficient.

4.4.2. *Effectiveness of the Moore machine strategies.* The analysis so far has confirmed Miller's premiss that challenging games elicit more complex behaviour but we have not examined the outcomes to evaluate the success of the Moore machine strategies. Miller has recorded the frequency of each outcome in the final rounds of the tournaments (see Table 3 so we can examine how well the strategies play. If we are willing to interpret ordinal values as actual payoffs, we can compute average outcomes for the players from the frequencies in Miller's repeated play encounters. These values are shown in Figure 10 as green crosses, together with red circles of the Nash equilibria. The performance of the strategies is quite effective as a glance at the figure shows.

- We assume 97% frequency of an outcome as a fixed solution strategy. The strategies are quite effective in the 36 games with dominant strategies for both players (top right quadrant of each layer), achieving the dominant strategy equilibrium solution in all games except three, 111 (Prisoner's Dilemma) with 24.6% frequency where the equilibrium solution is Pareto-dominated and 211 and 411 with 94%, where the equilibrium solution is not socially efficient. In all cases, both players face payoff pattern 1.
- In the 72 dominance solvable games (top left and bottom right quadrants), the equilibrium solution is achieved with 97% frequency in 56 games, including all with efficient, fair payoffs. Sixteen show more equivocal behaviour. In six of these games, 221 and 412, 231 and 413, 241 and 414, the equilibrium is Pareto-dominated and the more efficient outcome frequently results. In two games, 212 and 421, an inefficient equilibrium outcome is often avoided for the more efficient one. In six more games, 213 and 431, 214 and 441, 112 and 121, a fair, equally efficient outcome is sometimes selected over an unfair Nash equilibrium. In 252 and 425, an outcome with opposite payoffs for the players is sometimes reached. Of the sixteen games, only the last does not contain a payoff pattern of 1.
- There are 18 games with two Nash equilibria.
 - Nine are coordination games (including Stag Hunt) with a (4,4) outcome dominating another equilibrium. All these games are




FIGURE 10. Average outcomes for Moore machine players (green cross) compared to Nash equilibria (red circle)

successfully solved to achieve the dominant outcome except one, 333, where a reduced 94.3% frequency is achieved. Miller attributes this to an extreme penalty - payoff 1 - for trying to achieve the efficient outcome and failing to coordinate.

- The other nine games with two equilibria are battles-of-the-sexes with neither equilibrium dominating. Chicken is also in this set. These games are apparently harder to play. While equilibrium outcomes are almost always achieved (the exception is 122 Chicken where a fair and equally efficient outcome occurs once in nine) they are not efficient in 123 and 132, 124 and 142, the cases where the two equilibria are not equally efficient. In

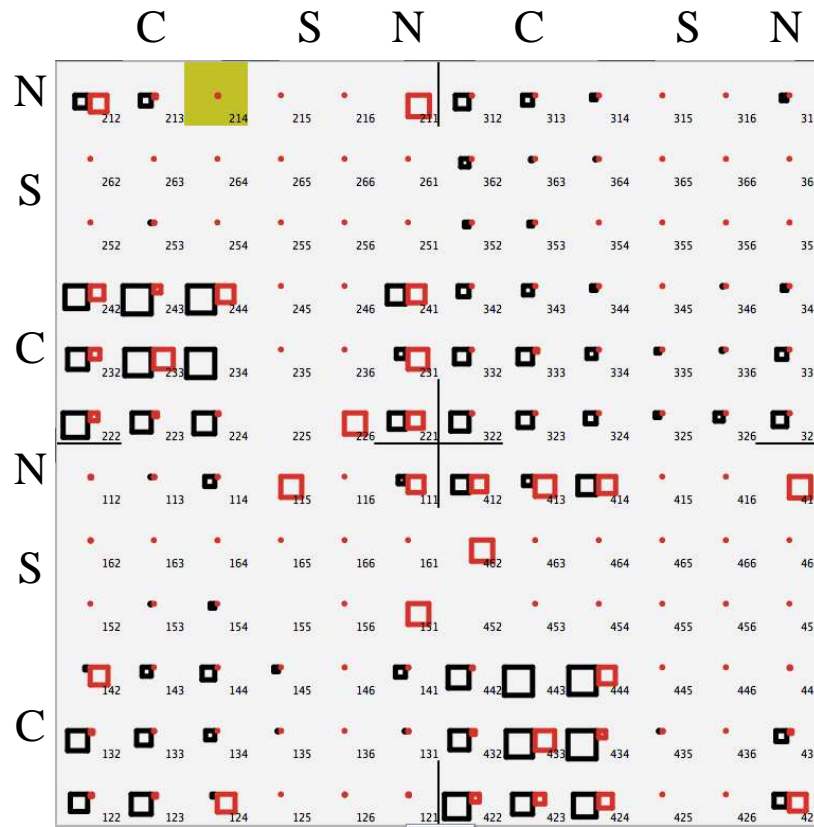


FIGURE 11. Efficiency and Complexity: Shortfall from most efficient outcome in red; average complexity of players in black

the four cases of equal efficiency, there is not fair distribution. In spite of the relative failure in these games compared to the coordination games, the complexity of the playing strategies is similar.

- The 18 games with no equilibria inspired the most complex strategies. That makes are the most difficult games but considerable success was achieved. In the ten games that had an efficient outcome of (4,3) or (3,3), this was found as the most frequent outcome in all cases. The other eight are games of pure conflict, including the constant sum games 234 and 443. These games had the most variation in outcome but the outcomes are reasonably efficient and reasonably fair, comparable to the battle-of-the-sexes games.

4.5. Complexity and Success. In Figure 11 efficiency and complexity are plotted together. Efficiency is calculated as the maximum total payoff available in the game for any output minus the actual total payoff for the Miller solution. It is shown as a red square with larger squares indicating a greater shortfall from maximum possible efficiency. Complexity is the total complexity for both players with only positive complexities shown as black squares, larger for more complex strategies. Simple strategies are not displayed. In general, the games where there is a single Nash equilibrium are simple to play with efficient outcomes. The notable exceptions are the seven, in an ‘L’ shape centred at game 111, the Prisoner’s Dilemma, where the dominant strategy outcome is Pareto-dominated. Games without a dominant strategy are more challenging and the Moore machines are more complex though reasonably successful at finding efficient solutions.

Thus, we can notice where strategies, simple and complex, have “failed to notice” opportunities for greater efficiency. When complex strategies have not achieved efficiency, a hypothesis could be that the machines did not possess enough potential for complexity to improve the result. These are the truly challenging games.

To examine the relationship of complexity and success more closely, another display in Figure 12 shows four values: success and complexity for both players. Again simple strategies are ignored; only complex strategies are shown. The Row player’s complexity, in black, is in the bottom left of the display, while Column’s is in the top right. Success is measured to a less rigorous standard: the shortfall from the payoff expected according to the Kalai-Smorodinski bargaining solution. The shortfall from the K-S value is recorded in red. A player’s payoff may exceed the KS value, in which case it is shown in blue. Row’s payoff success is shown at bottom right and Column’s at top left.

Computed from strict ordinal payoffs, the K-S values are always equal for both players so the diagram shows the inequity of payoffs. Most of the unfair distribution of payoffs occurs in simple games with dominant strategies: in many cases (layers 2 and 4), it is the player with the dominant strategy who receives less. By contrast, the payoffs are much more equitably distributed in the more challenging games with no dominant strategy patterns.

5. CONCLUSIONS

A search for games that present challenging situations to players should focus on player’s eye views of the play, not on the ultimate outcomes. This implies attending to preferences and to strategies for trying to deal with them. We have organized the ordinal 2×2 games according to preferences

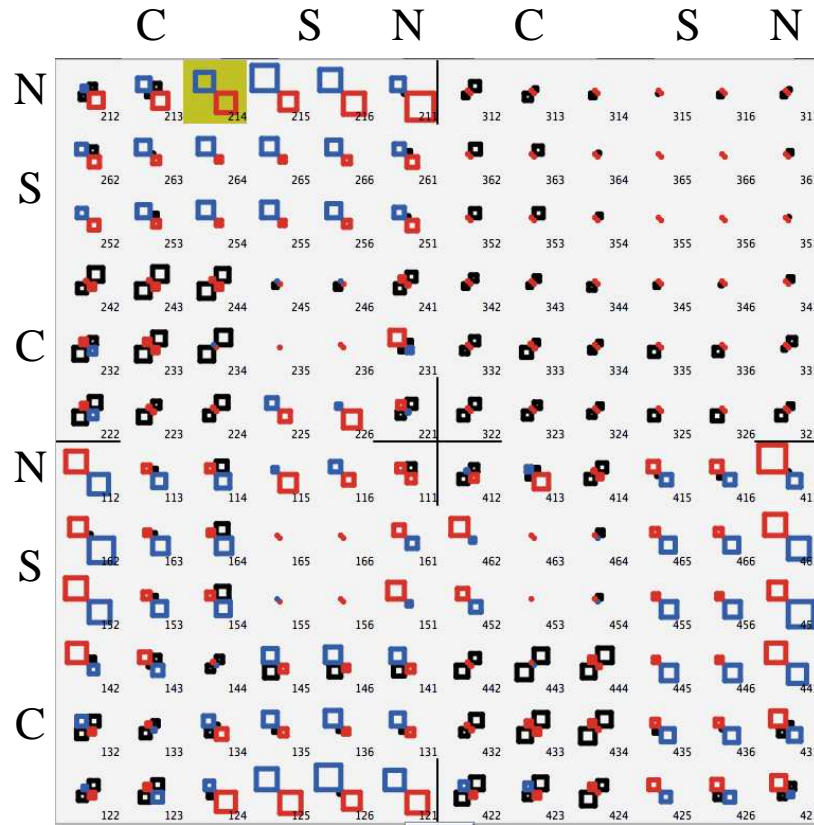


FIGURE 12. Player complexity and success.
 Lower left: Row complexity; lower right: Row payoff
 Upper right: Column complexity; upper left: Column payoff
 Red: deficit from KS bargaining payoff
 Blue: excess over KS bargaining payoff

and have found that the resulting structure clusters games with similar properties near each other. Miller’s experiments demonstrate that some games require more complex strategies to play successfully than others. The surprising result of displaying the complexity measures on the periodic table is that the primary determinant of a challenging game is the individual payoff pattern a player is faced with. Having a dominant strategy makes playing easier.

The games where the complex behaviour appears to work are the ones that we expect to be challenging because they do not admit of a simple strategy such as dominance solvability. The efficient and fair outcomes in the cyclic games with no Nash equilibrium are remarkable results of Miller’s experiments.

As Figure 12 shows, the strategy of dominance solvability, while simple, is not always successful. In the most blatant failures, the Pareto-dominated dominant-solvable outcomes of the Prisoner's Dilemma and its neighbouring games, the complex strategies have reacted and followed better strategies. (These are precisely the games where Miller observed temporal deviation in the long term proportional outcomes though we have not discussed this aspect of the behaviour in detail here.)

However, the Moore machines were taken in by dominance solvability in many other games where fairer and even more efficient solutions were possible. In the dominance solvable games 212(55), 213(56), 214(44), for example, Row has a dominant strategy that produces Nash equilibria solutions. The payoff distributions of these outcomes are unfair *against* Row, payoffs of 2 versus payoffs of 4 for Column. Row can do better by playing against the dominant choice. This is a subtler, presumably more complex, situation than the Pareto-dominated Nash equilibria. The Moore machines playing Row made some move toward better solutions: these games ended at the Nash equilibrium only 61.9%, 70.8% and 89.6% of the time, respectively, resulting in improved payoffs of 2.66, 2.28 and 2.08.

Perhaps more potential complexity is needed to capture the challenge of these games. In that case, our primary finding - that facing a dominant strategy makes for relatively easy play - could well be disproved.

6. ACKNOWLEDGMENT

Thanks to John Miller for sharing the data from his complexity experiments.

REFERENCES

- [1] Axelrod, Robert. "Effective Choice in the Prisoner's Dilemma" in *Journal of Conflict Resolution*, Vol. 24, No. 1: 3-25. 1980.
- [2] Axelrod, Robert. "More Effective Choice in the Prisoner's Dilemma" in *Journal of Conflict Resolution*, Vol. 24, No. 3: 379-403. 1980.
- [3] Daskalakis, Constantinos, P.W. Goldberg, and C.H. Papadimitriou. "The Complexity of Computing a Nash Equilibrium." *Communications of the ACM*, vol 52(2). 2009.
- [4] Kalai, Ehud and Meir Smorodinski. "Other Solutions to Nash's Bargaining Problem" in *Econometrica*, Vol. 43: 513-518. 1975.
- [5] Miller, John. *A Strategic Taxonomy of Repeated 2 x 2 Games Played by Adaptive Agents* Draft document. Carnegie Mellon and Santa Fe Institute. 1991.
- [6] Miller, John. *Complex Adaptive Systems*. Princeton NJ. Princeton University Press. 2007.
- [7] Nash, John. "The Bargaining Problem". *Econometrica* 18(2): 155-162. 1950.
- [8] Nash, John. "Noncooperative Games" *Ann. Math.* 54. 1951.
- [9] Robinson, David and David Goforth. *The Topology of The 2 x 2 games : A New Periodic Table*. London UK. Routledge. 2005.

Current address: Department of Mathematics and Computer Science
Laurentian University

E-mail address: dgoforth@cs.laurentian.ca

URL: <http://www.cs.laurentian.ca/dgoforth/home.html>

Current address: Department of Economics

Laurentian University

E-mail address: drobinson@laurentian.ca

URL: <http://www.laurentian.ca/drobinson/home.html>


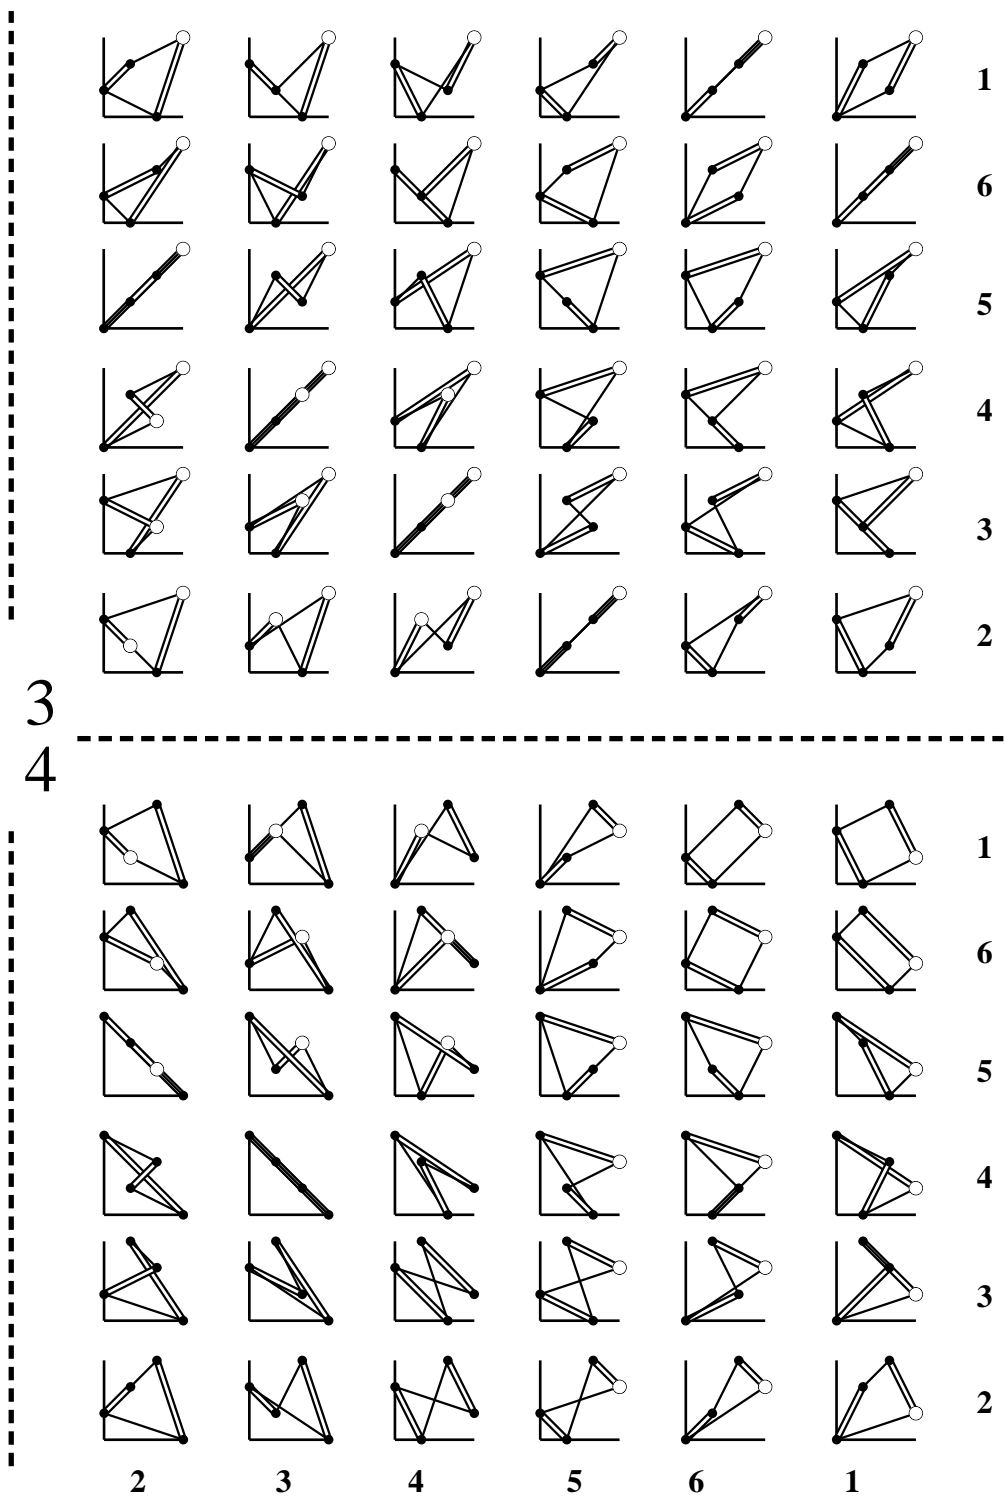


FIGURE 13. The Periodic table of the 2×2 games



Game index:layer, row, column

1	44/33/22/11	.	a	.ND	P.R	P.C	P..	99.6	0.2	0.2	0.0	-0.5	0.4	-0.1	-0.50	-0.55
2	44/33/12/21	.	a	.ND	P.R	P.C	P..	99.5	0.2	0.2	0.0	-0.4	0.1	-0.3	-0.58	-0.63
3	44/32/23/11	S	a	.ND	P.R	P.C	P..	99.5	0.2	0.2	0.0	-1.0	-0.7	-1.6	-0.59	-0.63
4	44/32/13/21	.	a	.ND	P.R	P.C	P..	99.5	0.3	0.1	0.0	-0.3	-0.8	-1.1	-0.50	-0.54
5	44/31/13/22	S	a	.ND	P.R	P.C	P..	99.6	0.2	0.2	0.0	-0.8	0.5	-1.4	-0.58	-0.63
12	22/41/14/33	S	c	PND	..R	..C	...	24.6	12.7	13.8	48.9	0.9	-0.1	0.8	5.81	5.45
55	24/43/11/32	.	j	.NR	..R	P..	P..	61.9	32.8	2.3	3.0	0.7	0.7	1.4	1.15	1.21
60	44/21/12/33	S	k	.N.	P..	P..	PN.	99.3	0.3	0.3	0.1	0.7	-0.2	0.5	-0.58	-0.62
63	44/12/21/33	S	k	.N.	P..	P..	PN.	94.3	0.2	0.3	5.2	1.1	0.5	1.6	-0.50	-0.54
78	22/31/43/14	.	n	P..	P..	4.0	4.4	81.4	10.2	1.4	0.8	2.3	0.33	0.39
G#	Payoffs	S	RG	0,0	0,1	1,0	1,1	0,0	0,1	1,0	1,1	rcpx	ccpx	scpx	maxTD	aggTD

Column Legend

G#	RG number
Payoffs	payoff matrix
S	(S) symmetric or (.) not
RG	Rapoport and Guyer classification (e.g., a: two dominant strategies, no conflict)
0,0 0,1 1,0 1,1	(P or .) Pareto-dominated; (N or .) pure Nash equilibrium; (D,R,C,.) both, row, col or no dominant strategy
0,0 0,1 1,0 1,1	frequencies of observed outcomes
rcpx, ccpx, scpx	normalized complexity for row, column and aggregate
maxTD, aggTD	normalized maximum outcome frequency temporal deviation, aggregate deviation for all outcomes