

Is Knowledge Power? Information in Games

Ariana Strandburg-Peshkin
Santa Fe Institute REU Program 2008
Mentors: John Miller and Willemien Kets

ABSTRACT

Information asymmetries are commonplace in strategic interactions. In this paper, we explore how having a small amount of information about an opponent's strategy can affect play and ultimately payoffs in games. Using an agent-based model in which agents play one another in iterated two-player games and evolve their strategies, we study what happens when we give some agents knowledge of their opponents' strategy lengths and leave others ignorant. We find that agents with information that adapt their strategies to exploit this information can achieve higher payoffs against certain populations whose agents do not evolve, but do no better against populations without information that are themselves evolving. The reason is that when populations coevolve, agents without information can adapt their strategies so as to nullify any possible advantage the agents with information may have. Moreover, agents with information are at a disadvantage in games with two differently-preferred equilibria, such as the battle of the sexes and chicken. We present an explanation based on the stability of different equilibria.

INTRODUCTION

The outcomes of strategic interactions between individuals often depend on the information that these individuals have. This information can be asymmetric, in that one individual may have more information than the other. An interesting type of information difference occurs when one agent has some information about its opponent while the other does not. In this paper, we explore the effect of this type of information in complex adaptive social systems. These systems are composed of populations of interacting agents that evolve their interaction strategies over time (Holland and Miller, 1991). Examples of this type of system include markets, ecosystems, political systems, and social networks. Using an agent-based model and evolutionary algorithm, we explore how giving some agents information about the potential complexity of their opponents' strategies can affect their strategies, the outcomes of games they play, and the overall dynamics of a system composed of both "informed" and "uninformed" agents.

We choose to give agents the information of their opponents' potential for complexity for two reasons. First, we expect that this information may be relevant to an agent's strategy. For example, if an agent knows that an opponent has a very simple strategy and will play the same action in every game, that agent should play the best response to that action. However, if an agent knows that an opponent's strategy is potentially more complex, it may be beneficial to engage in more complicated interactions with that individual, such as reciprocating its moves. Human experiments have provided evidence that individuals may alter their strategies in games based on how they perceive their opponents (Eichberger, 2006). The second reason to choose potential for complexity as the information provided to an agent is

that it can be simply quantified via an agent's strategy length, as discussed below.

In this paper, we discuss how giving certain agents information about their opponents' potential for strategic complexity can affect strategic interactions in two settings. First we explore how information affects interactions when agents are playing against populations of agents whose strategies do not evolve. Next, we investigate the role of information when agents play against populations of agents whose strategies are evolving over time. We discuss situations in which information may be useful, when it is neutralized, and even when it may be harmful.

MODEL

Model Overview

We use an agent-based model in which agents play iterated 2x2 games (Fig. 1). The model uses a simple evolutionary algorithm in which agents evolve and compete with one another to reproduce. In each generation of the model, agents play one another in 2x2 iterated games and receive payoffs based on the outcomes of these games. Some agents are then selected to reproduce and are placed into the next round (agents with higher payoffs are more likely to reproduce). Finally, with a certain probability, agents' strategies are mutated. This process of game play, selection, and mutation (Fig. 2) is repeated for a given number of generations, thus allowing the system to explore different strategies over time.

	Action A	Action B
Action A	$u_1(A,A), u_2(A,A)$	$u_1(A,B), u_2(A,B)$
Action B	$u_1(B,A), u_2(B,A)$	$u_1(B,B), u_2(B,B)$

Fig. 1 – Generalized 2x2 Payoff Matrix. Player 1 is considered to be the row player, while player 2 is the column player. Their payoffs for each combination of actions are represented by u_1 and u_2 respectively.

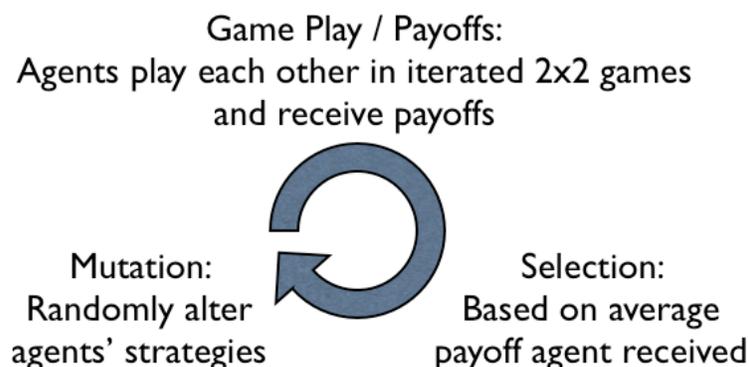


Fig. 2 – Model Schematic

Agents

Agents are modeled as finite automata (specifically, Moore machines). Finite automata consist of a series of S evolvable states. Each state has an action and transition function. The action is what the agent will play in the current round of the game. The transition function specifies which state to move to based on the opponent's last play (Fig. 3). In the case of 2x2 games, studied here, the transition function contains two possible transitions, which specify which state to move to if the opponent plays action A and which state to move to if the opponent plays action B. Finite automata generally begin each round of games in a specified starting state, although informed agents may begin in one of two starting states, as we discuss below. If agents do not receive additional information, the assignment of actions to states and the transition functions define a strategy. As we describe below, some agents receive additional information prior to play. In that case, the definition of a strategy includes an additional variable, which we specify below. Mutations of actions and transition functions allow agents' strategies to evolve over time. The detailed mechanism for mutation is discussed below.

The finite automata framework allows agents to condition their actions on the previous actions of their opponents. Finite automata are a preferable model for agent behavior because they provide inherent flexibility and realistic and adjustable bounds on the computational abilities of individual agents. For these reasons, finite automata have been used in prior studies of games in complex adaptive social systems (e.g. Miller 1996, Miller 2002, Miller 2003).

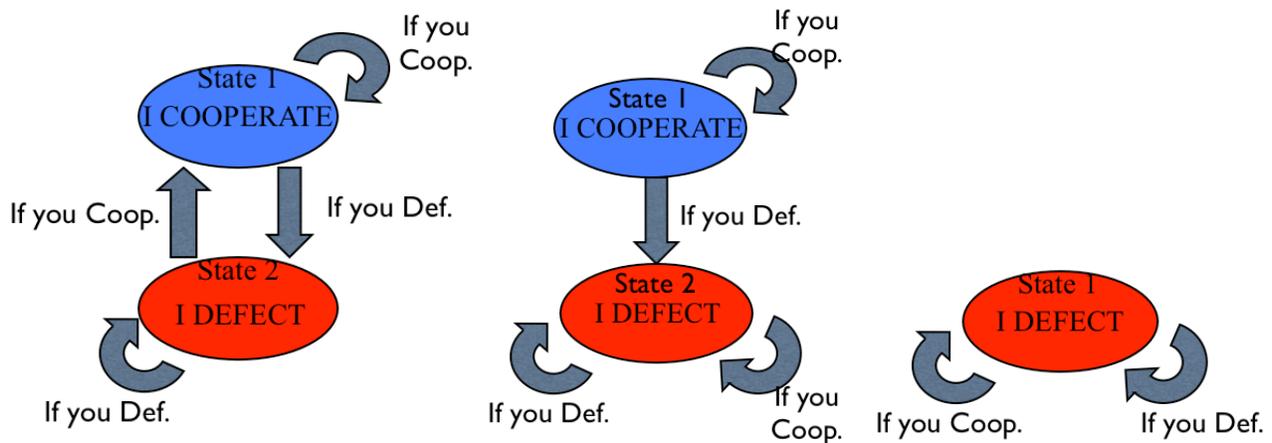


Fig. 3 – Schematic Diagrams of Finite Automata. Action choices in this case are “cooperate” and “defect.” The action each state plays is written inside the oval, and its transitions are represented by arrows. The left panel depicts the strategy known as Tit-For-Tat. Tit-For-Tat starts by playing “cooperate” and then reciprocates its opponents previous move. The center panel shows the so-called “trigger” strategy. This strategy begins by playing “cooperate” and then defects forever if its opponent ever defects. The right panel shows the strategy known as “all-defect.”. This strategy plays “defect” in every round regardless of its opponent's actions.

Strategy Length

The strategy length of an agent, l , is defined as the minimum number of states necessary to fully describe its strategy. For example, a strategy which always plays action A could be fully described using only one state, which plays action A and always transitions back to itself. More complicated strategies require the use of more states. Tit-For-Tat, for example, requires two states, and a strategy

which fully specifies actions based on an opponent's plays in the last n rounds could require up to 2^n states (Miller, 1996). Thus, an agent's strategy length is a measure of the potential complexity of its strategy. Note that an agent's strategy length does not necessarily provide an accurate measure of the *realized* complexity of its strategy, only its *potential*. For example, a random strategy might be minimally described by many states, while the actual complexity of a random strategy is very low. To find the minimal description of agents' strategies, we use the algorithm outlined in Harrison (1966, p. 311).

Informed Agents

In order to explore the role of information, we introduce modified agents which can receive and process information. These agents have the same basic structure as the agents described above, with one critical alteration. Uninformed agents (which receive no information prior to game play), always start in the same state, which we denote state 1. Informed agents (which *do* receive information prior to game play) condition their starting state on the information they receive. In this case, informed agents receive the information of their opponent's (minimized) strategy length, and may start in state 1 or state 2 depending on this information. Informed agents process the strategy length information via a cutoff value, c , which is an adjustable number between 0 and the maximum number of states, S . If the strategy length of an informed agent's opponent is at or below the informed agent's cutoff, the agent will start in state 1. Otherwise it will start in state 2 (Fig. 4). In this way, an informed agent is able to condition its strategy on the length of its opponent's strategy (and, therefore, on its potential for complexity). Hence, for informed agents, a strategy specifies the assignment of actions to states, a transition function, and a cutoff value. In addition to evolving the actions and transition functions of states, informed agents also evolve their cutoffs.

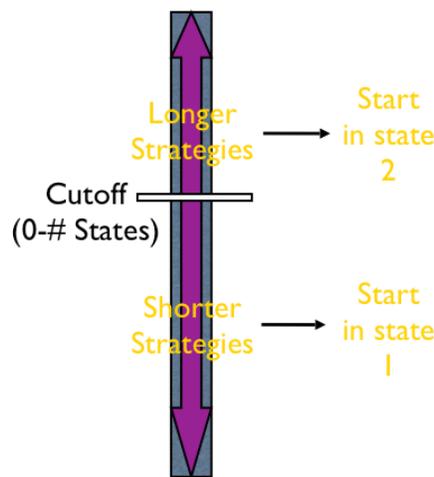


Fig. 4 – The cutoff mechanism of informed agents

Split Populations

We split the population of agents into two subpopulations, referred to as subpopulation 1 and subpopulation 2. Games are played between agents from different subpopulations, while selection operates within a subpopulation. In other words, agents in subpopulation 1 interact via games with agents in subpopulation 2 (and vice versa) during the game play phase of a generation. However, during the selection phase, agents in subpopulation 1 only compete against other agents in subpopulation 1 to

reproduce, and similarly for subpopulation 2 (Fig. 5). This split population framework was briefly discussed in Miller, 1996 (p. 109).

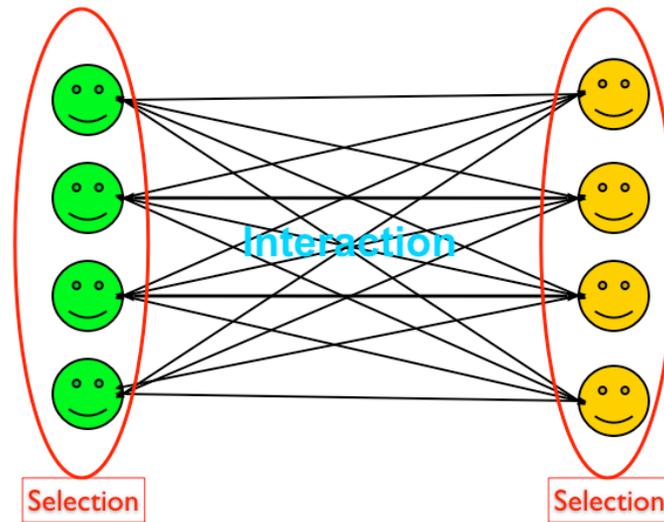


Fig. 5 – Split Population Framework

The rationale behind the split population framework is two-fold. From a practical standpoint, having split populations avoids the problem of interactions between informed agents. Since informed agents can start in two states, depending on their opponent's strategy length, they do not have a well-defined strategy length. Hence, they cannot interact via games with one another. Splitting the population into a subpopulation of informed agents and a subpopulation of uninformed agents solves this problem. Split populations also have the added benefit of allowing for asymmetric games, in which one subpopulation acts as the row player and the other acts as the column player. From a theoretical standpoint, split populations model a situation in which individuals interact with members of another group but compete with individuals within their own group. The situations of sellers and buyers in a market or species in different niches which interact with one another in an ecosystem are good examples of where split populations occur naturally.

Simulations

We seed the model with two subpopulations of N agents. For the reasons discussed in the previous section, these subpopulations may both be composed of uninformed agents, or one may be composed of uninformed agents and the other of informed agents. In most cases, these initial subpopulations will contain agents playing random strategies. A random strategy is generated by choosing actions and transitions for each state uniformly at random from the set of possible actions and the set of all states, respectively. In the case of informed agents, the cutoff is set to a random value between 0 and S , the maximum number of states. In each generation, agents play one another in G iterations of a 2x2 game and receive payoffs. Selection and mutation then occur, and the process repeats.

Game Play

During game play, every agent in subpopulation 1 plays every agent in subpopulation 2 (and vice versa) in iterated 2x2 games, consisting of G rounds. Agents receive payoffs depending on the outcomes of these games, and their payoffs are summed throughout the generation.

Selection

The selection method used here is tournament selection. In this type of selection, two agents are picked uniformly at random from the subpopulation (with replacement). The agent with the higher payoff reproduces and is placed into the next generation. This process is repeated until each new subpopulation contains N agents. In this way, the size of each subpopulation remains constant, while agents with higher payoffs are more likely to reproduce. Note that this selection method does not guarantee that the best strategy in the population will be kept or that the worst will be eliminated.

Mutation

After selection, agents' strategies in the new population are mutated. Agents are selected uniformly at random for mutation with a probability M . If an agent is chosen for mutation, one of its states is chosen uniformly at random to be altered. Half of the time, this state's action is altered and half of the time one of its transitions (chosen uniformly at random) is altered. If the mutating agent is an informed agent, its cutoff is mutated with a probability C .

The selection and mutation mechanisms outlined above were previously used in Miller, 2002 and Miller, 2003.

Unless otherwise noted, the parameters used in simulations were:

$$N = 25$$

$$S = 12$$

$$M = .5$$

$$C = .25$$

$$G = 10$$

RESULTS

Information in Static Populations

Static populations are populations in which only one subpopulation is allowed to evolve, while the other remains fixed. We study the effect of information in static populations in the context of the Prisoner's Dilemma. The Prisoner's Dilemma payoff matrix is shown in Figure 6.

	Cooperate	Defect
Cooperate	2 , 2	0 , 3
Defect	3 , 0	1 , 1*

Fig. 6 – Prisoner's Dilemma Payoff Matrix. The asterisk denotes the unique Nash equilibrium.

We allow subpopulations of informed or uninformed agents to play against various fixed subpopulations of uninformed agents. Agents in the latter subpopulation do not undergo selection or mutation, so their strategies remain constant over time. We refer to this subpopulation as the fixed subpopulation. This allows us to explore the effect of strategy length information without the further complication of changing information.

We find that when a subpopulation of evolving agents plays against a subpopulation of agents whose strategies are fixed, strategy length information can sometimes be helpful in distinguishing between

agents that play simple strategies and agents that could potentially play more complex strategies. For example, subpopulations of informed agents do better than subpopulations of uninformed agents against a fixed subpopulation of uninformed agents, half of which always play defect (all-D, Fig. 3, right), and half of which play Tit-For-Tat (TFT, Fig 3, left). We ran 50 simulations of 500 iterations each and found that informed agents received a higher average payoff than uninformed agents. The difference, although small, was statistically significant ($p=.000002$). The mechanism behind the informed agents' advantage is simple. Because all-D is a strategy of length 1 (it takes only one state to fully describe it) and TFT is a strategy of length 2, informed agents are able to discriminate immediately between the two strategies by setting their cutoffs to 1. They are then able to defect with the all-D agents (below their cutoffs) and cooperate with the TFT agents (above their cutoffs).

This advantage is smaller than it may seem, however, because uninformed agents are able to distinguish between the two strategies quickly as well. For example, evolving uninformed agents playing against the TFT/all-D fixed population might evolve a TFT-like strategy, or a “trigger” strategy which starts by cooperating and then defects forever if its opponent ever defects (Fig. 3, center). Informed agents using a cutoff only outperform uninformed agents using TFT or a trigger strategy in the first round of the game, thus the apparent advantage is greatly reduced, which explains the small size of the difference in average payoff.

Even this small informed agent advantage disappears when agents are faced with less contrived subpopulations. Obviously, informed agents fare no better against a subpopulation of agents whose strategy lengths are all the same. In addition, against a fixed subpopulation playing iterated Prisoner's Dilemma with random strategies, informed agents do no better than uninformed agents. It is easy to see why this is – in a random subpopulation, reciprocal strategies are unlikely, so the optimal strategy is to defect, regardless of an opponent's strategy length. More generally, random subpopulations do not contain much information, so informed agents are not advantaged.

Information in Coevolving Populations

When informed subpopulations and uninformed subpopulations are both allowed to evolve, the informed advantage again disappears. Simulations of all 256 symmetric 2x2 games, as well as the 78 canonical games (Rapoport and Guyer, 1978) were run. To judge the effect of information when subpopulations coevolve, we looked at the average payoff of agents in each subpopulation for each game over 500 iterations. In no game did the informed agents outperform the uninformed agents. The average payoffs of the informed agents were not statistically different from the average payoffs of the uninformed agents, with some notable exceptions, discussed below.

At first glance, this result may suggest that strategy length information is simply not useful, but our previous work shows that strategy length information can sometimes be useful when agents play against fixed subpopulations. The disappearance of the information advantage in coevolving populations suggests that one of four explanations may be responsible. First, it may be that the uninformed subpopulations are too random for information to be of significant use. Just as information is not useful against fixed subpopulations of agents with random strategies, perhaps agents' strategies in a coevolving world are so noisy that they approximate randomness. On the other hand, it may be that coevolving subpopulations are too homogeneous. As discussed above, strategy length information is rendered useless if all agents in a subpopulation have the same strategy length. However, by looking at the strategy length distributions over time, we observe at least some heterogeneity. Third, it may be that uninformed subpopulations evolve too fast so that informed agents are not able to configure their

cutoffs in time. This claim makes intuitive sense, as informed agents need to adapt their cutoff value as well as the assignment of actions to states and the transition functions, whereas uninformed agents only need to adjust the latter two. To test this possibility, subpopulations of informed and uninformed agents were allowed to evolve against a fixed subpopulation of uninformed agents with random strategies. We observed the number of generations it took for the payoffs to reach a steady state. The tests were inconclusive, however. The time to reach a steady state is indeed longer for informed agents, but the difference is not statistically significant.

Lastly, it is likely that uninformed agents may evolve their strategy lengths, and may use this ability to place themselves on the side of the informed agents' cutoffs that is most preferable to them. We tested the feasibility of this last explanation by alternating which population evolved within a given number of iterations. The results showed that uninformed agents were able to evolve their strategy lengths to move to the more advantageous side of a cutoff. For example, assume we again start with a fixed subpopulation of uninformed agents in the iterated Prisoner's Dilemma, half of which play TFT and half of which play all-D. We allow an initially random subpopulation of informed agents to play against this fixed population and evolve. Just as before, most agents in the informed subpopulation evolve a cutoff of 1, above which they cooperate and below which they defect. If we then fix the informed subpopulation and allow the uninformed subpopulation to play against it and evolve, most uninformed agents evolve longer strategies that mostly play defect, thus taking advantage of the informed agents' willingness to cooperate with agents whose strategy lengths are greater than one. A simple example of such a strategy is shown in Fig. 7. This simple experiment shows that uninformed agents have the ability to evolve their strategy lengths to take advantage of differences in informed agents' actions on either side of their cutoffs.

Thus, while strategy length information may be beneficial to agents in certain fixed settings, the advantage that information yields is easily neutralized when uninformed agents are allowed to evolve their strategies.

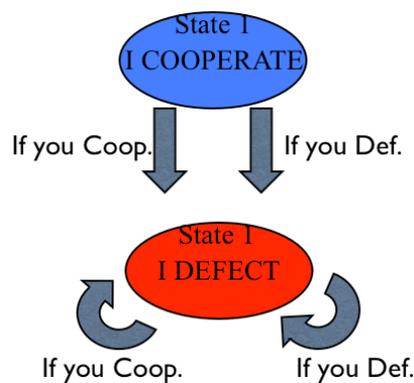


Fig. 7 – A mostly defecting strategy of length 2. Such strategies might evolve in the uninformed subpopulation if agents if the informed subpopulation had evolved to cooperate with agents whose strategy lengths were above 1 and defect otherwise.

Information as a Disadvantage

Standard game theory assumes that people behave as Bayesian updaters, thus information is weakly beneficial because one can always ignore it. However, when we relax this assumption, the result that information cannot be harmful no longer applies. Indeed, we have found that in certain games, the

agents that have strategy length information receive lower payoffs on average than the uninformed agents with whom they are interacting. These “information-disadvantage” games all have a similar payoff structure. They are coordination or anti-coordination games with two equilibria in which the row player prefers one equilibrium whereas the column player prefers the other. Common examples of this type of game include chicken and the battle of the sexes (BoS), shown in Figures 8 and 9. BoS was extensively studied here.

	I	U
I	1, 3*	0, 0
U	0, 0	3, 1*

Fig. 8 - Battle of the sexes payoff matrix. Asterisks indicate the pure Nash equilibria. Informed agents prefer the equilibrium in which both agents play I whereas uninformed agents prefer the equilibrium in which both play U.

	A	B
A	2, 2	3, 1*
B	1, 3*	0, 0

Fig. 9 - Chicken payoff matrix. Asterisks indicate the pure Nash equilibria. Both players prefer the equilibrium in which they play action A and their opponent plays action B.

Dynamics of BoS

The dynamics of coevolving populations of agents playing BoS are rather complex. In general, the system tends to reach what we will refer to as “temporary steady states” in which players almost always coordinate on one of the two equilibria (with a low level of misses). In many cases, the agents coordinate on the same equilibrium during each round of their interactions, although sometimes the agents alternate which equilibrium they play in various patterns. These temporary steady states are punctuated by rapid transitions between them. (Fig. 10). We will call a temporary steady state in which agents almost always coordinate on the U strategy (which is preferred by the uninformed agents) an “uninformed-preferred temporary steady state” (UPTSS) and one in which agents almost always coordinate on the I strategy (which is preferred by the informed agents) an “informed-preferred temporary steady state” (IPTSS).

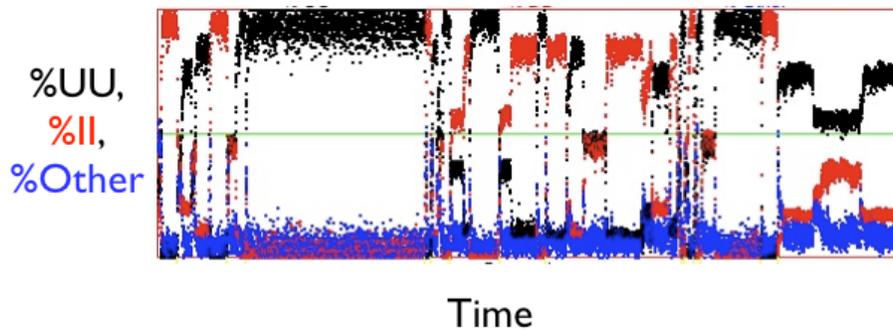


Fig. 10 – Dynamics of BoS. The graph shows the fraction of the time the two equilibria (II and UU, black and red respectively) were played over various generations. Blue represents the fraction of the time the agents did not coordinate.

Information Disadvantage in BoS

In BoS and similar games, when a subpopulation of uninformed agents is allowed to evolve and interact with an evolving subpopulation of informed agents, the average payoff of the informed subpopulation is found to be lower than the average payoff of the uninformed population. 190 runs of 500 iterations each were run. The average payoff of the informed agents was 1.67 points per game, whereas the average payoff for uninformed agents was 1.93 per game. These results were statistically significant ($p < .000001$).

There are two possible explanations for the discrepancy in payoffs between the informed and uninformed agents. First, it is possible that the informed agents are worse at staying in their preferred temporary steady state (the “defense” hypothesis). Alternatively, this could be viewed as the uninformed agents being better at knocking the informed agents out of their preferred temporary steady state than vice versa. If the defense hypothesis is correct, we should observe the average length of time (number of generations) spent in each UPTSS to be greater than the average length of time spent in each IPTSS. Second, it is possible that the informed agents are worse at getting into their preferred temporary steady state once the system is no longer in a temporary steady state (the “offense” hypothesis). If the offense hypothesis is correct, we should observe the frequency of UPTSS's to be greater than the frequency of IPTSS's. In other words, the system should enter a UPTSS more often than an IPTSS.

In order to test these hypotheses, we look at the frequency and average duration of UPTSS's and IPTSS's in 3 runs of 10,000 iterations each (we use longer runs here to minimize the effect of temporary steady states being cut off by a run ending). We find that within each run there are on average 30 UPTSS's and 15 IPTSS's. On average the UPTSS's last 176 iterations and the IPTSS's last only 99 iterations. These results provide support for both the offense and the defense hypotheses, suggesting that both mechanisms are at play.

Informed Disadvantage and Number of States

In order to test the effect of the computing capacity of agents on the disadvantage of informed agents observed in BoS, we altered the total number of states, S , that each agent has. Preliminary results suggest that the payoff difference between informed and uninformed agents decreases as more states are added (although the result that informed agents do worse than uninformed agents holds true for all values of S tested). Figure 11 shows the results of this experiment.

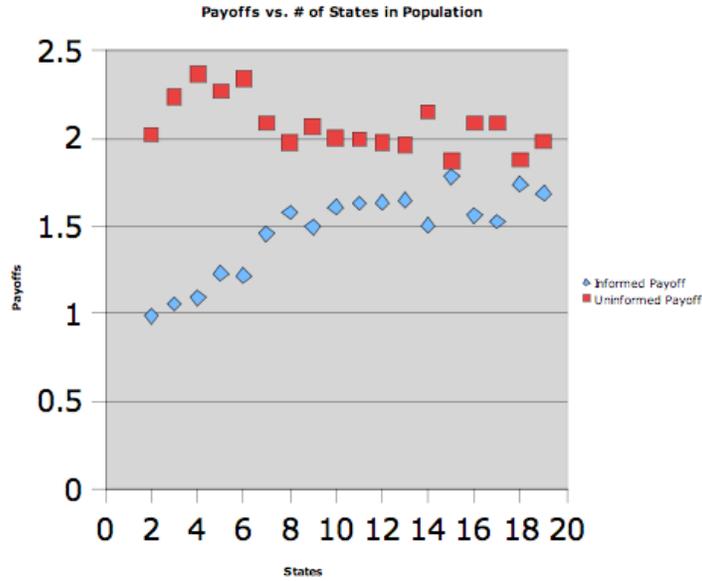


Fig. 11 – Effect of Number of States on Payoff in BoS

These results suggest that the information disadvantage in BoS is robust to changes in the number of states that agents may use. In addition, the decreasing trend in the data suggest that the mechanism behind the information disadvantage is at its strongest when the total number of states is small. Thus, in order to explore this mechanism, it is beneficial to look at a simplified case of BoS dynamics.

Simplified Case

In order to uncover the mechanism behind the information disadvantage, we studied a vastly simplified case in which agents have only 2 states and play one-shot BoS games (each agent plays a single game of BoS with each other agent). This simplified case displays similar qualitative results to more complex cases involving more states and larger number of iterated games. Therefore, analyzing this case in depth may help us shed some light on the mechanism behind the information disadvantage in BoS in the general case.

In this simplified case, there are only four distinct strategies possible in each population:

Uninformed Population Strategies	Informed Population Strategies
u_1 – play U , strategy length = 1	UU – always play U
u_2 – play U , strategy length = 2	II – always play I
i_1 – play I , strategy length = 1	IU – play I if opponent’s strategy length = 1, play U if opponent’s strategy length = 2
i_2 – play I , strategy length = 2	UI – play U if opponent’s strategy length = 1, play I if opponent’s strategy length = 2

The payoffs each agent receives are dependent on the fraction of opponents playing the four strategies. The payoffs π_u and π_i for the uninformed and informed agents respectively are shown below, where f_j

represents the fraction of a given subpopulation playing strategy j and $\Delta_I=(f_{UU},f_{IU},f_{UI},f_{II})$ and $\Delta_U=(f_{u1},f_{u2},f_{i1},f_{i2})$ denote the population state faced by the informed and uninformed agents, respectively.

Uninformed Agent Payoffs	Informed Agent Payoffs
$\pi_u(u1, \Delta_I) = 3f_{UU} + 3f_{UI}$	$\pi_i(UU, \Delta_U) = f_{u1} + f_{u2}$
$\pi_u(u2, \Delta_I) = 3f_{UU} + 3f_{IU}$	$\pi_i(II, \Delta_U) = 3f_{i1} + 3f_{i2}$
$\pi_u(i1, \Delta_I) = f_{II} + f_{IU}$	$\pi_i(IU, \Delta_U) = 3f_{i1} + f_{u2}$
$\pi_u(i2, \Delta_I) = f_{II} + f_{UI}$	$\pi_i(UI, \Delta_U) = f_{u1} + 3f_{i2}$

Importantly, each type of agent's payoffs depends on the fractions of two types of strategies in the other population. We find that the genetic drift (random mutations in the absence of selective pressure) between these two strategy types is the driving force behind the information disadvantage.

To gain insight into the information disadvantage, we compare the stability of the IPTSS and the UPTSS.

Stability of the Informed-Preferred Temporary Steady State and the Uninformed-Preferred Temporary Steady State in the Simplified Case

Assume that we start in the IPTSS, with the informed subpopulation being primarily composed of II agents and the uninformed subpopulation being composed of both i_1 and i_2 agents. Because the payoffs for i_1 and i_2 agents are equal, genetic drift between these two strategies may occur. This may lead to an asymmetry in the sizes of the i_1 and i_2 populations. Due to this genetic drift, at times the subpopulation of uninformed agents may be composed almost exclusively of either i_1 or i_2 agents. This occurs because once a slight asymmetry in frequencies of strategy types arises, it can be magnified in the following generations due to the fact that the more frequent strategy is more likely to be selected for reproduction. Once either i_1 or i_2 agents make up most of the uninformed subpopulation, the informed agents' strategies may drift to what we will call a split strategy, UI or IU. Given a sufficient number of one of these split strategies, it is possible for u_2 or u_1 uninformed agents to successfully invade the uninformed population and displace the other strategies because the payoffs for u_1 or u_2 agents are higher than the i_2 or i_1 payoffs (due to the prevalence of UI or IU agents respectively). The u_1 or u_2 agents then grow in number, and this has the further effect of advantaging the split strategy agents over the II agents. Thus, a feedback loop between u_1 and UI agents (or u_2 and IU agents) is created and the system breaks out of the IPTSS.

One may think that a similar argument could be made for the UPTSS, but this is not the case. This is because the feedback mechanism always advantages both the split strategy and the u strategy, thus it always favors the UPTSS. To see this, consider a system in the UPTSS. Due to genetic drift, the uninformed subpopulation becomes composed of mostly u_1 agents, and therefore there is genetic drift in the informed subpopulation towards the UI strategy. At this point, i_2 agents have the opportunity to invade, but because the uninformed agents prefer the UU equilibrium, the u_1 agents consistently outperform the i_2 agents. Therefore, the system stays in the UPTSS.

The dynamics of the simple case scenario can be visualized in the form of the flow chart in Figure 12.

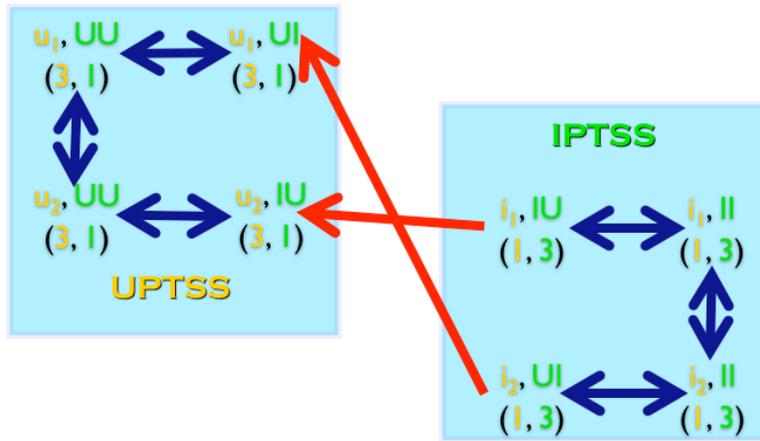


Fig. 12 – Flow diagram for the dynamics of BoS in the simplified case. The blue double-headed arrows represent states of the population between which there may be genetic drift because the payoffs are equal for the player whose strategy changes when moving between population states. The red arrows represent selective pressures that will push the system in one direction. From the diagram one can see how the UPTSS is more stable than the IPTSS because the selective pressure arrows both point towards the UPTSS.

DISCUSSION AND CONCLUSION

Results

We studied how information affects agents' payoffs in several cases. When playing the Prisoner's Dilemma game against certain subpopulations of agents whose strategies are fixed, agents with information on their opponents' strategy lengths outperform agents without information. This result is not surprising, as one would expect that additional information about an opponent would be an advantage in at least some cases. The extent of this information advantage is limited, however, as against more complicated (i.e. less contrived) subpopulations, the value of the strategy length information decreases, and against random or completely homogeneous populations this information is worthless.

However, when subpopulations of agents are allowed to coevolve, it was found that information is not an advantage in any game. This counterintuitive result is probably due to an "arms race" in which informed agents continually evolve their cutoffs to distinguish between different types of uninformed agents, and uninformed agents continually evolve their strategy lengths to be on the more advantageous side of the cutoff. Thus, any advantage that agents may derive from information can be neutralized when agents coevolve.

Third, in games with two differently-preferred equilibria such as battle of the sexes and chicken, information was found to be a disadvantage in that agents that receive information receive lower payoffs on average than agents that do not receive information. This result runs contrary to the standard intuition that additional information can never hurt because it can always be ignored. In a simple case of battle of the sexes, we showed that in using the information for temporary gain, informed agents may contribute to their own demise by creating a feedback loop which moves the system into a

state which is less advantageous to them. This process depends critically on genetic drift between strategies that receive equal payoffs in both subpopulations.

The result that information can be harmful in some cases raises questions about the assumption in game theory that players are Bayesian updaters. If players are Bayesian updaters, having additional information cannot be harmful because it can always be ignored. By contrast, we have shown here that if this assumption is relaxed, having some additional information about an opponent may be harmful in some cases.

It is of note that in the case of battle of the sexes and similar games, information sometimes benefits individuals within a subpopulation (those with split strategies – UI or IU) but is detrimental to the subpopulation as a whole. In addition, the dynamics of battle of the sexes also highlight the difference between short-term and long-term benefit, as informed agents with split strategies might do better or equally well in the short term, but their existence causes their subpopulation's long-term disadvantage.

Future Directions

In this paper we have explored the effect that information on an opponent's potential for complexity can have on interactions in evolving settings. We have shown how this information can be beneficial, useless, or even harmful to the agents that receive it. These results are fairly limited, however, in that they relate to only one type of information and only one method of information processing. Future work could explore the effect of different types of information on strategic interactions in evolving settings. One possible type of information about an agent's strategy that might be explored is the extent to which that agent reciprocates. It is likely that this type of information could help sustain cooperation in games like the Prisoner's Dilemma, and may also have interesting effects in games like the battle of the sexes. Another type of information that could be explored is some measure of an agent's actions in previous games. This type of information represents an individual's reputation.

In addition to investigating the effects of other types of information, other methods of information processing should also be explored. The cutoff mechanism of the informed agents is clearly an oversimplified model of how individuals might use information on their opponents, and one can imagine more elaborate methods such as multiple cutoffs, which would allow agents to more precisely classify their opponents based on the information they receive.

In our model, we provide agents with exact knowledge of their opponents' strategy lengths. However, in real situations, individuals most likely never receive such precise information. The addition of noise in information transmission would make the model more realistic and would test its robustness.

Finally, in our current model there is no cost associated with having more complex strategies. This is clearly inaccurate, as more complex strategies require more computing power on the part of agents, and computing power is costly. The addition of a cost based on an agent's strategy length would introduce a bias in the direction of simplicity of strategy, and could lead to different results. In particular, the genetic drift mechanism may no longer operate in battle of the sexes, thus it could be that the information disadvantage would be reduced or even eliminated.

WORKS CITED

- Eichberger, Jürgen, Kelsey, David and Schipper, Burkhard C., "Granny versus Game Theorist: Ambiguity in Experimental Games." (September 26, 2006).
- Harrison, Michael A.. Introduction to Switching and Automata Theory. New York: McGraw-Hill, 1966.
- Holland, John, and John Miller. "Artificial Adaptive Agents in Economic Theory." The American Economic Review 81.2 (1991): 365-370.
- Miller, John H., Carter T. Butts, and David Rode. "Communication and Cooperation." Journal of Economic Behavior and Organization 47 (2000): 179-195.
- Miller, John H., and Scott E. Page. Complex Adaptive Systems: An Introduction to Computational Models of Social Life (Princeton Studies in Complexity). Princeton: Princeton University Press, 2007.
- Miller, John H., and Scott Moser. "Communication and coordination." Complexity 9.5 (2004): 31-40.
- Miller, John H.. "The Coevolution of Automata in the Repeated Prisoner's Dilemma." Journal of Economic Behavior and Organization 29 (1996): 97-112.
- Rapoport, Anatol, and Melvin Guyer. "A Taxonomy of 2x2 Games." General Systems XXIII (1978): 125-136.