Title
Balancing Fairness and Efficiency in Repeated Societal Interaction

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN
1069-7977

Authors
Ben-Asher, Noam
Lebiere, Chrisitan
Oltramari, Alessandro
et al.

Publication Date
2013

Peer reviewed

Balancing Fairness and Efficiency in Repeated Societal Interaction

Noam Ben-Asher (noamba@cmu.edu)[1]
Christian Lebiere (cl@cmu.edu)[2]
Alessandro Oltramari (aoltrama@andrew.cmu.edu)[2]
Cleotilde Gonzalez (coty@cmu.edu) [1]

[1]Dynamic Decision Making Laboratory, Department of Social and Decision Sciences, Carnegie Mellon University
Pittsburgh, PA 15213 USA

[2]Department of Psychology, Carnegie Mellon University
Pittsburgh, PA 15213 USA

Abstract

Fairness and efficiency are important aspects that influence cooperation in social dilemmas. During a repeated interaction, they have the potential to serve as competing goals for the decision maker. The ability to balance between fairness and efficiency depends, among other things, on available information regarding mutual accountability for the outcomes in an interaction. In this paper, we examine how information regarding mutual interdependencies influences the interplay between fairness and efficiency in repeated Chicken Game. We distinguish between three possible types of fair behavior: mutual cooperation, alternating cooperation, and mutual destruction. Our results show that the first two types of fairness are positively correlated with the availability of social information. In contrast, mutual destructive fairness is not sensitive to the availability of information and is generally avoided. We also find that without information regarding mutual interdependencies, unfairness increases in parallel with efficiency. When social information is available, however, increases in fairness is coupled with a decrease in efficiency, and the best compromise between fairness and efficiency is reached when mutual interdependencies are learned through repeated experiences. We highlight the significance of our results for fair and efficient interaction in repeated social interactions.

Keywords: Efficiency; Fairness; Cooperation; Game Theory; Interdependence; Information; Social Interaction

Introduction and Background

In our daily lives, we constantly face situations in which our well-being and success depends on the actions of others. Whether the interactions occur between individuals or between organizations, there is mutual accountability for the outcomes. For example, the interaction between two toddlers that learn to share a toy by taking turns holds some similarities with the interaction between companies that adopt a competitive brinkmanship pricing policy while trying to gain control over a certain market. If the toddlers are not willing to behave in a fair manner they will both end up screaming and none of them will play with the desired toy. An alternative behavior is that both toddlers decide simultaneously to switch interest to other toys and thus eliminate the conflict. A more mutually beneficial behavior is where the toddlers will share the toy so one can play with it for a while and then the other will have the joy of playing with the desired toy as well. However, if only one of the toddlers gets to play with the toy and the other does not get the same opportunity, feelings of unfairness and frustration that might lead to aggressive behavior when facing similar conflicts in the future can arise.

Such social conflicts are well captured in Game Theory using the Chicken Game (CG), as introduced by Russell (1959). According to this game, two drivers are heading towards each other on a single lane road from opposite directions at full speed. Just before colliding, each of the drivers has to choose simultaneously and independently between driving straight towards a possible collision (i.e., Dare) or turning the steering wheel (i.e., Swerve) and avoiding the accident. As represented in the game's payoff matrix (see Table 1), this is a prototypical dangerous game, because a player has to risk the lowest payoff [-10] to have a chance of winning the highest payoff [10]. Under reasonable assumptions for single-trial CG, the best outcome is for a player that Dares while the other player Swerves [10,-1]; the second-best for each is if both Swerve [1,1]; the third-best is for a player that Swerves while the other player Dares [-1,10]; and the worst for each is if both Dare [-10,-10], because then the outcome is mutually destructive. Thus, the best strategy in single-trial CG depends on the opponent‘s expected behavior and a player can maximize the outcome by doing the opposite of what the other player does. However, for repeated CG (infinitely repeated in theory; finitely repeated with unknown endpoint in practice), successful alternations, where one player wins the highest payoff in one round and then the other player wins the highest payoff in the next round, is the best strategy to obtain a joint maximum outcome. This type of cooperation corresponds well to the situation where, over time, the well-being of one player depends of the well-being of the other.

The need to consider the well-being of the other challenges traditional economic theories, which assume that people act selfishly. In contrast, there is growing experimental evidence that actual behavior is also shaped by factors inconsistent with pure selfishness (e.g., Dufwenberg & Kirchsteiger, 2004; Fehr & Schmidt, 1999; Roth, 1995). The psychology literature also provides evidence for other factors that individuals consider, beyond their own well-being, and acknowledges the importance of reciprocity by incorporating it into models of human behavior. Heider (1958) introduced the idea that causal inference, where one takes into account another person’s motives and situational constraints, as an important cognitive process for perceiving social contexts. Similarly, Game Theory incorporates these ideas by considering altruism and fairness (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999). Altruism can be simply defined as an interaction in which people care not only about their own well-being but also about others. This over-simplified definition is extended through fairness into different directions by incorporating distributive concerns (Bolton & Ockenfels, 2000), inequity aversion (Fehr & Schmidt, 1999), and reciprocity theories (Rabin, 1993).

CG is particularly suitable to studying different aspects of fairness in social interaction. In its basic form, fairness is achieved when one player's outcome is identical to the other player's. Considering the payoff matrix presented in Table 1, fairness can happen when both players make the same decision in a given round. It is possible that both players selected Dare (i.e., mutual Dare) and therefore received an outcome of [-10, -10], or when both players selected Swerve (i.e., mutual Swerve) and received the outcome of [1, 1]. Repeated CG adds another kind of fair interaction where the players alternate for consecutive rounds, [10, -1] followed by [-1, 10] and so on, resulting in alternating cooperation. This type of fairness is also the optimal strategy in repeated CG (the one that results in the highest long-term outcome). It is possible that fair coordination like [1, 1] would be easier to achieve than alternating coordination because the latter require a more complex coordination of choices. However, once a state of alternating coordination is achieved, it might be more stable compared to simple coordination (Rapoport, Guyer, & Gordon, 1976).

One common but unrealistic assumption in research on strategic social interaction is that players possess full information about their interdependence, usually presented using a payoff matrix. However, several studies demonstrated the value of certain types of information in well-known social games, such as the repeated Prisoner's Dilemma (Camerer, 2003; Rapoport & Chammah, 1965). Recently, Gonzalez and Martin (2011) proposed the Hierarchy of Social Information (HSI), a theoretical framework for conceptualizing and organizing the major categories of interpersonal information that may play a role in social interactions. Martin, Gonzalez, Juvina, and Lebiere (2012) used this framework to examine the effects of information on cooperation in repeated Prisoner's Dilemma. Their findings reveal a generally positive impact of information on joint performance and satisfaction. They showed that an increase in cooperation with an increase in the availability of social information was driven in part by players’ greater willingness to reciprocate the other player's prior cooperation, and concluded that players possessing more interdependence information were more likely to enforce social norms of reward and retribution. Such social norms depend heavily on perceptions of fairness and how it is expressed in repeated interaction. Furthermore, the availability of interdependence information can influence awareness of fair and unfair aspects of the interaction. Thus, it is possible to assume that having enough information to compare between one’s own and another's payoffs might be a minimal requirement for fairness. However, it is currently unclear how more information systematically influences fairness, and more specifically the impact of different types of fairness that can occur as part of the interaction between two interdependent players. Furthermore, there is a need to understand the costs of maintaining fairness during a social interaction. Some social settings emphasize the tradeoff between fairness and performance. For example, in one-shot Ultimatum Game, a proposer maintains fairness and cooperation with a responder by decreasing her own personal gain. Other social settings completely disentangle such relationship between fairness and performance like in one-shot Dictator Game, where there is no need to maintain fairness and the proposer can keep all the gains for herself. Forsythe and colleagues (1994) demonstrated that on average offers in the Ultimatum game were higher than in the Dictator Game, confirming the tradeoff between efficiency of cooperation and fairness. However, people do care about fairness and in some situations are willing to sacrifice some of their own to maintain fairness (Bolton & Ockenfels, 2000, Güth, Kliemt, & Ockenfels). Thus, it is possible that a fair distribution of payoffs is appealing from a social perspective, but it might have some negative influence on performance in repeated interaction.

Table 1: Chicken game payoff matrix, with Action A denoting Dare and Action B denoting Swerve. The cells show a pair of outcomes (x, y) where x is the payoff to Player 1 and y is the payoff to Player 2.

		Player 2 Action	
		A(Dare)	B (Swerve)
Player 1 Action	A (Dare)	-10, -10	10, -1
	B (Swerve)	-1, 10	1, 1

The current paper presents an experiment to examine how information regarding mutual interdependencies can influence fairness in CG. We start by providing background information on a repeated CG game that was used to collect the behavioral data, and present the four levels of information proposed in the HSI framework (Gonzalez & Martin, 2011). Then, we distinguish between three types of fair outcomes in repeated CG and examine the interactions between fairness and the availability of social information. Following that, we focus on the relations between fairness and efficiency in alternating cooperation. Finally, we discuss how awareness of interdependence may encourage efficient and fair behavior in real-world interactions and describe potential future directions of this research.

Experiment

Participants and Procedure

Participants (N=240; 120 pairs, 45% of whom were women, M_{age}=22.8, SD_{age}=4.53) were recruited to a computer laboratory at Carnegie Mellon University, and randomly paired to play 200 rounds of repeated CG over the laboratory's network. Players were not told the number of rounds in the game, and the rounds were not numbered in the course of play. In each round, the two anonymous members of a pair (seated in different rooms without having met one another) chose simultaneously between buttons labeled Action A and Action B with payoffs as in Table 1. These payoffs were converted to incentive pay (one cent per point) beyond $10 base pay. One pair was excluded from the analysis due to communication failure during the experiment.

Information Conditions

All participants saw their own action and payoff in each round. Thirty pairs (60 participants) were assigned to each of the four conditions that determined the amount of information available about their interdependence. These conditions were modeled after the layers of the Hierarchy of Social Information (HSI) outlined by Gonzalez and Martin (2011). We briefly describe each of the information conditions here.

In the “Individual” condition, players were not informed that they were interacting with another player, so the selection of an action in each round was most likely perceived as an independent binary choice between two options with probabilistic payoffs. Participants may have realized that the outcome probabilities for each of the two actions were not static, as they in fact varied with the other player’s actions, but this could more easily be attributed to a computerized process that shifted exogenously or in response to their own actions.

In the “Minimal” condition, players knew that their outcomes depended on the actions of another player and vice versa, yet they still did not know the other’s specific actions and payoffs. With this information, individuals may have been able to speculate about the other’s motivations, but it would remain difficult to infer the other’s actions and payoffs.

Next, pairs in the “Experiential” condition saw each other's actions and outcomes in each round. This information allowed players to reason about the mutual interdependencies through repeated experiences.

Finally, in the “Descriptive” condition, in addition to seeing each other’s actions and outcomes, players were shown the complete payoff matrix (as in Table 1) from the outset and throughout the repeated interaction.

Results

We analyzed the three different fair outcomes described above: mutual Swerve, where both players receive a small positive payoff of +1 with the risk of experiencing a moderate loss of -1 if the other player deviates from fairness; mutual Dare where both players lose 10 points with the potential of winning +10 if the other player deviates; and the alternating cooperation, which requires that both players successfully alternate between Swerve and Dare in a way that the payoffs in one round are [+10, -1] and players then get [-1, +10] in the following round.

General Fairness Preferences

Overall, the joint proportion of fair interactions increased with more information given: .38 in the Individual condition, .50 in the Minimal condition, .71 in the Experiential condition, and .74 in the Descriptive condition. However, as illustrated in Figure 1, behavior varied for the different types of fairness. Next, we examine each of the fair interactions separately.

Alternating Cooperation

The average proportion of alternating cooperation varied significantly across the four information conditions and generally increased with greater availability of information, $F(3, 115)=7.086$, $p<.001$. As shown in

Figure 1, the average proportion of alternating cooperation in the Individual condition was extremely low (M=.03, SD=.05). Knowing that the payoffs depend on the decisions of another human player provided in the Minimal condition (M=.1, SD=.19) increased the alternating cooperation significantly, $t(57)=2.106$, $p=.040$. The average proportion of alternating cooperation increased substantially more in the Experiential condition (M=.29, SD=.34), where each player observed the actions and payoffs of the other player, and was significantly higher than the Individual and Minimal conditions, $t(58)=4.202$, $p<.001$ and $t(57)=2.575$, $p=.012$, respectively. Similarly, in the Descriptive condition (M=.22, SD=.28) the average proportion was significantly higher than the Individual condition and marginally higher than the Minimal condition, $t(58)=3.765$, $p<.001$ and $t(57)=1.880$, $p=.065$, respectively. No significant differences were found between the Descriptive and Experiential conditions, $t(58)=.847$, p=ns.

Mutual Swerve

The average proportion of mutual Swerve varied significantly across the four information conditions, $F(3,115)=2.78$, $p=.044$. As shown in Figure 1, on average, the proportion of fair rounds where both players swerved increased with the availability of information: M=.21 (SD=.23) in the Individual condition, M=.21 (SD=.17) in the Minimal condition, M=.27 (SD=.24) in the Experiential condition, and M=0.36 (SD=.27) in the Descriptive condition. The average proportion of mutual Swerve in the Descriptive condition was significantly higher than in the Individual and Minimal conditions, $t(58)=2.341$, $p=.023$ and $t(57)=2.466$, $p=.017$, respectively. These results suggest that participants in the Descriptive condition tended to prefer the less risky option that could yield a small gain but also a small loss compare to the risky option that could yield a large gain or a large loss. Also, the ability to see the interdependencies as they are represented by the payoff matrix highlights the fairness of mutual Swerve. Thus, it is possible that some synergy between these two interpretations of the social information can explain the increased proportion of mutual Swerve.

Mutual Dare

When both players select Dare in the same round, they both receive the same negative payoff. As seen in Figure 1, the average proportion of mutual Dare was not affected by the availability of information, $F(3,115)=.37$, p=ns. The average proportion of rounds where both players dared remained similar while the available information increased: M=.15 (SD=.08) in the Individual condition, M=.18 (SD=.10) in the Minimal condition, M=.15 (SD=.11) in the Experiential condition, and M=.16 (SD=.21) in the Descriptive contrition. The SD in the Descriptive condition is relatively higher compared to the other conditions, mainly due to two pairs in the Descriptive condition who mutually dared in more than 80% of the 200 rounds.

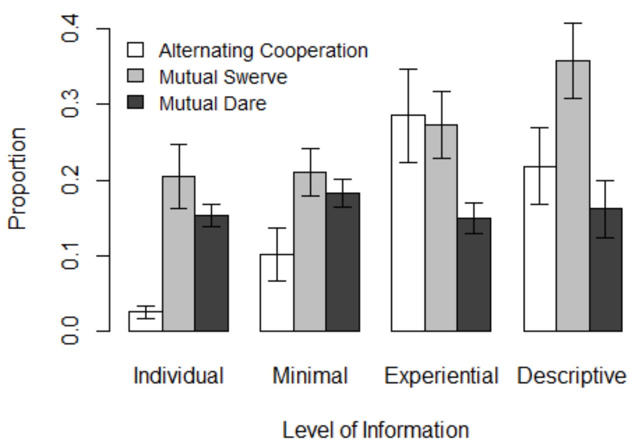

Figure 1. Average proportion of fair outcomes for four levels of social information across 200 rounds.

Efficiency and Fairness tradeoff in Alternating Cooperation

Rapoport et al. (1976) suggested the alternating cooperation index (k) as a measure suitable for evaluating cooperation between two players in games where the optimal collective strategy is coordinated alternations. This measure uses the frequencies of the asymmetric payoffs in the following form: k= (DS) + (SD) - |(DS) - (SD)|

Where DS and SD refer to the number of times that the asymmetric payoffs occurred. The maximal value of k is achieved when DS and SD occur exclusively and with equal frequencies, which would correspond to perfect alternation.

Considering the setting of the study presented here, the maximal value of k is 200 and it is achieved when DS=SD=100. The minimal value of k is 0 and it represents a game in which one player dominated the other player throughout the 200 rounds, resulting in DS=200 or SD=200.

The sum of DS and SD reflects the exclusiveness of the asymmetric payoffs compared to other possible payoffs. Thus, it represents the efficiency of the alternating cooperation. On the other hand, the absolute value of the difference between DS and SD reflects the balance or fairness between the two players and serves as a penalty for unfair behavior. The magnitude of the penalty determines whether one player dominated the other. For the sake of simplicity and clarity, we decompose k to its terms (i.e., efficiency and unfairness penalty) and converted them to proportions by dividing each term by the number of rounds (i.e., 200). Furthermore, we use 1 – the proportion of unfairness penalty to calculate the proportion of fairness.

Similar to alternating cooperation, k varied significantly across the information conditions. We observed a general increase of k with more information: M=.25 (SD=.18) in the Individual condition, M=.32 (SD=.23) in the Minimal condition, M=.46 (SD=.28) in the Experiential condition, and M=.39 (SD=.28) in the Descriptive condition.

To gain a better understanding of how the availability of social information influenced the alternating cooperation index (k) and especially the relations between its efficiency and fairness components, we analyzed each of the terms that construct k separately. As seen in Figure 2, fairness increased significantly with the availability of social information, $F(3,115)=12.778$, $p<.001$. The average fairness in the Descriptive condition (M=.91, SD=.11) was significantly higher than the average fairness in the Individual (M=.61, SD=.26) and Minimal (M=.72, SD=.28) conditions, $t(58)=5.826$, $p<.001$ and $t(57)=3.601$, $p<.001$, respectively. Similarly, the average fairness in the Experiential condition (M=.88, SD=.18) was significantly higher than in the Individual and Minimal conditions, $t(58)=4.817$, $p<.001$ and $t(57)=2.869$, $p=.005$.

As shown in Figure 2, the analysis of the efficiency term also indicated that efficiency varied significantly across the information conditions, $F(3,115)=3.030$, $p=.032$. However, we found an trend opposite of fairness, where efficiency in the Descriptive condition (M=.48, SD=.25) was significantly lower compared to the Individual (M=.64, SD=.21) and Minimal (M=.61, SD=.17) conditions, $t(58)=-2.756$, $p=.008$, $t(57)=2.296$, $p=.025$.

The analysis above indicates that the availability of social information influenced the alternating cooperation index in general and differentially influenced the efficiency and fairness terms that compose it. It seems that the relatively high fairness in the Experiential and Descriptive came at the cost of decreased efficiency. However, the decrease in efficiency when moving from the Minimal condition to the Experiential condition is not significant, while the increase in fairness is. This suggests that best compromise between fairness and efficiency was reached in the Experiential condition.

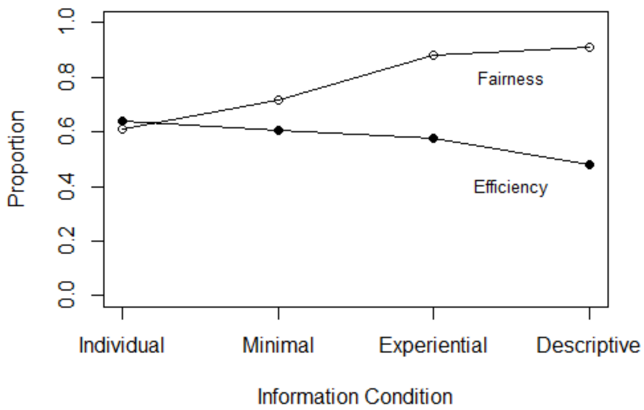

Figure 2: Average proportion of efficient and fair alternating cooperation for four levels of social information across 200 rounds.

To gather a better understanding of the relations between fairness and efficiency, we analyze efficiency and fairness at the pair level, for each of the information conditions (see Figure 3). Results indicate that efficiency and fairness were strongly and negatively correlated in the Individual and Minimal conditions, r=-.73, p<.001 and r=-.56, p=.002, respectively, in contrast to the Experiential and Descriptive conditions, where no significant correlations were found, r=.10, p=ns. and r=-.02, p=ns. This finding suggests that as the efficiency of alternating cooperation increased in the Individual and Minimal conditions, fairness between the two players decreased. This means that one player dominated the other more often, resulting in an unfair and heedless behavior.

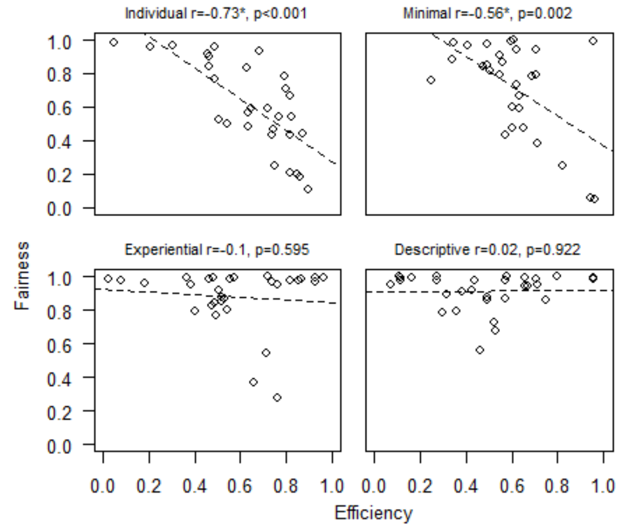

Figure 3: Correlations between proportions of efficiency and fairness in alternating cooperation for each pair, for four levels of social information, across 200 rounds.

Discussion and Conclusion

We tested how fairness is influenced by the availability of social information regarding mutual accountability in a social dilemma, as represented by the Chicken Game. Drawing upon Gonzalez and Martin's (2011) Hierarchy of Social Information, we find that fairness in such social dilemmas depended on the availability of information and in general increased when more information was available to the decision makers. Thus, information availability moderates unfairness and can increase fair social behaviors.

First, the overall high proportion of fair outcomes in the Descriptive and Experiential conditions implies that the availability of detailed information regarding interdependencies with others elicit fair behavior. This holds mainly for the two constructive interactions (mutual cooperation and alternating cooperation), which leads to positive outcomes for both sides of the conflict. On the other hand, when there is no awareness of the counterpart's conditions, unfairness increases on the expense of fairness. This results in an antisocial interaction where one exploits the other, even unintentionally and without being aware of the nature of the interaction.

Mutual destructive fairness stands out from the other two types of fairness as it was not influenced by the availability of information. Even without the knowledge of interdependence, players managed to avoid the mutually destructive escalation of the conflict which leads to negative outcomes for both sides. However, individuals possessing interdependence information were more likely to behave in a social manner and preferred other fair interactions. In contrast, individuals that did not possess such information were more likely to exhibit unfair and unsocial behavior where one exploits the other.

We also find that when only experiential information is available, the proportions of mutual cooperation and alternating cooperation are relatively similar. In contrast, when descriptive information is also provided, there is a greater preference towards the fair outcome resulting from mutual cooperation, compared to alternating cooperation. This might relate to the increased complexity of coordination that alternating cooperation requires, compared to mutual cooperation (Rapoport et al., 1976). Alternatively, it is possible that players concluded from the descriptive information (i.e., the payoff matrix) that there is a relatively low risk in choosing to Swerve compared to Dare and thus preferred this option more.

The alternating cooperation index (k) provided us with important insights regarding the relations between fairness and efficiency within this type of interaction. Both fairness and efficiency were influenced by the availability of social information. While fairness increased with more information regarding mutual accountability, efficiency decreased. Moreover, increases in fairness between the information conditions were steeper and more drastic, compared with decreases in efficiency. This finding suggests that under certain conditions, increases in fairness might have some cost, and lead to a minor decrease in the interaction's efficiency. This finding in repeated interaction is consistent with the tradeoff between efficiency and fairness in one-shot Ultimatum Game (Forsythe et al., 1994; Güth et al., 2003). However, it seems that for repeated interactions, fairness is more sensitive to the changes in the availability of social information compared to efficiency. The best compromise between fairness and efficiency was achieved when mutual interdependencies were learned only through repeated experiences. In this condition, fairness increased significantly while any decrease in efficiency was insignificant. This suggests that the availability of complete information, as in the Descriptive condition, do elicit fairness, but not necessarily the most efficient kind of fairness.

Examining the relations between fairness and efficiency at the pair level provided supporting evidence for these ideas. Where there is no or limited information regarding mutual accountability (i.e., the Individual and Minimal conditions), we find a decrease in fairness as the efficiency of a pair increases. This stands in contrast to the Experiential and Descriptive conditions, where we find no correlation between fairness and efficiency at the pair level. This demonstrates how overall the availability of information serves as a guard for fairness, and the efficiency of the interaction within the boundaries of fair behavior, depends on the specific interaction between individuals.

A key implication of this study is the importance of information to the fair resolution of conflicts. We show that social fairness is sensitive to the availability of information regarding mutual interdependencies between the members of the community. It is illuminating to see that the mere knowledge of interdependence with another person given in the Minimal condition is insufficient to promote alternating cooperation compared to the Individual condition. Further availability of social information increases cooperation and fairness. Thus, a preliminary requirement of fair conflict resolution should involve sharing information that sheds light on the interdependencies between the different entities involved in the conflict.

Even though it seems that overall fair behaviors in the Descriptive and Experiential conditions were somehow similar and less sensitive to the different way in which social information was conveyed, studies of individual decision making distinguish between these two sources of information (descriptive and experiential), and demonstrate how decisions from experience and description differ (e.g., Hertwig & Erev, 2009). Recently, a similar distinction has been made in social dilemmas involving more than one decision maker (Martin et al., 2012). We believe that descriptive and experiential information influences the way fair behavior evolves over time. Thus, further analysis should carefully examine and compare the dynamics of fair interaction over time for different levels of social information.

Acknowledgments

This research was supported by a Defense Threat Reduction Agency (DTRA) grant to Cleotilde Gonzalez and Christian Lebiere (HDTRA1-09-1-0053); and by the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. We are also thankful for the help and insightful comments of Hau-yu Wong.

References

Bolton, G., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, *90*(1), 166–193.

Camerer, C.F. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.

Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, *47*(2), 268–298.

Fehr, E., & Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817–868.

Forsythe, R., Horowitz, J. L., Savin, N. E., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior*, *6*(3), 347-369.

Gonzalez, C., & Martin, J. M. (2011). Scaling up instance-based learning theory to account for social interactions. *Negotiation and Conflict Management Research*, *4*(2), 110-128.

Güth, W., Kliemt, H., & Ockenfels, A. (2003). Fairness versus efficiency: An experimental study of (mutual) gift giving. *Journal of Economic Behavior & Organization*, *50*(4), 465-475.

Heider, F., (1958). *The psychology of interpersonal relations*. New York City: Wiley.

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences, 13*(12), 517-523.

Martin, J. M., Gonzalez, C., Juvina I., & Lebiere, C. (2012). A description-experience gap in social interactions: Information about interdependence and its effects on cooperation. Manuscript under review.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, *83*(5), 1281-1302.

Rapoport, A., & Chammah, A. M. (1965). *Prisoner's dilemma: A study in conflict and cooperation*. Ann Arbor, MI: University of Michigan Press.

Rapoport, A., Guyer, M. J., & Gordon, D. G. (1976). *The 2 x 2 game*. Ann Arbor, MI: University of Michigan Press.

Roth, A. E., (1995). Bargaining experiments. In J. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics*. Princeton, NJ: Princeton University Press.

Russell, B., (1959). *Common sense and nuclear warfare*. New York City: Simon & Schuster.