

Towards a Repeated Negotiating Agent that Treats People Individually: Cooperation, Social Value Orientation, & Machiavellianism

Full Paper

<Authors Redacted for Review>

ABSTRACT

We present the results of a study in which humans negotiate with computerized agents employing varied tactics over a repeated number of economic ultimatum games. We report that certain agents are highly effective against particular classes of humans: several individual difference measures for the human participant are shown to be critical in determining which agents will be successful. Asking for favors works when playing with pro-social people but backfires with more selfish individuals. Further, making poor offers invites punishment from Machiavellian individuals. These factors may be learned once and applied over repeated negotiations, which means user modeling techniques that can detect these differences accurately will be more successful than those that don't. Our work additionally shows that a significant benefit of co-operation is also present in repeated games—after sufficient interaction. These results have deep significance to agent designers who wish to design agents that are effective in negotiating with a broad swath of real human opponents. Further-more, it demonstrates the effectiveness of techniques which can reason about negotiation over time.

KEYWORDS

Human-Agent Negotiation; Individual Differences; Empirical Studies; Human-Computer Interaction

1 INTRODUCTION & BACKGROUND

Success in social tasks such as negotiation requires multi-faceted intelligence. The best human negotiators are not only able to formulate and deploy complex strategies, but also are able to take measure of their opponent and adjust tactics on the fly. Determining who to trust, when to be aggressive, and what words to speak in negotiation can have as much (or more) of an impact on outcomes as does the actual offers being made. Accordingly, any truly effective machine intelligence that would attempt to negotiate with a human would need to deftly handle these various negotiation aspects as well.

Agents that interact with humans should not simply adopt static, one-size-fits-all approaches. However, individual differences and human-agent relationships that evolve over time are difficult to understand intuitively. There are perhaps few places where these dynamics are more at the fore than in human-agent negotiation. This work examines the human-agent negotiation domain and highlights the importance of creating agents that incorporate both temporally-based models of

negotiation as well as models of individual user differences. We present the results of a user study in which four negotiating agents created using a collection of agent design “best principles”, are matched against human participants. The results indicate that while each agent can do well in a vacuum of user information, future agents that could automatically detect individual differences between human opponents will perform at a much higher level.

Even in a fairly constrained domain such as negotiation where the task of a given interaction is normally quite clear, human partners often act in unpredictable ways. Classical results from human-human literature have showed the use of anger to gain concession [20], and previous work has replicated these results using artificial agents [10]. Many other similar studies have been conducted which allow a general corpus of best design practices to be constructed. However, these classical models do not often take individual differences into account when outlining strategies for succeeding in negotiation, nor do they make continuous distinctions based on the number of interactions and the length of a given set of negotiations. Furthermore, they tend to ignore the obviously harmful effect of competitive or negative emotional techniques on trust and attitude toward the opponent. This may be a reasonable assumption on simple, one-shot games, but does not serve for more realistic, repeated interactions.

This work uses virtual agents to examine individual differences between human negotiators and study effects of agent strategy at different negotiation lengths. This allows us to create superior agents that utilize more complex models of interaction. This work represents another step in outlining several negotiation features that are important for the design of realistic agents. Namely, it explores the importance of acceptance rates over time as well as detecting individual differences. When designing agents that negotiate with humans, customizing their behavior to match their partner is not simply a good principle: *it provides statistical benefit*.

2 BACKGROUND & RELATED WORK

2.1 Why Negotiation?

Among the myriad uses for human-aware agents, negotiation is a critical subdomain. Indeed, negotiation is both a valuable end unto itself, as well as a stepping stone onto other human-agent interactions. A broad swath of human interaction can be thought of as falling within the umbrella of negotiation—from the obviously formal negotiation present when negotiating for a job or

buying a used car to the more flexible back-and-forth associated with deciding group dinner plans or conversational flow. Other sub-disciplines of human-computer interaction have connections to negotiation as well: preference elicitation is a critical component of many negotiating agents, and turn-taking strategies are essential for any natural language-based system.

It is thus very important that human-aware agents are designed that have good models of their human partners. Whether these agents are used directly for negotiation, where they either try to outdo their human partners or to teach them, or if they are used merely for their models of human behavior, the agents must be designed according to data-backed principles.

2.2 The Ultimatum Game

In an ultimatum game, one party is given a set of resources that have a certain amount of utility. That party (the proposer) is then allowed to offer the other party (the receiver) any number of the items. Once an offer is made, it cannot be changed. The receiver may then decide to accept or reject the offer. If the offer is rejected, all resources are lost for both sides. If the offer is accepted, then the resources are distributed according to the offer. Therefore, the dominant strategy (in the simple, non-repeated form game) is to accept any offer made in order to maximize points. However, in practice, most humans experiencing such a task will reject offers considered to be too unfair [8,14,17]. In the repeated-form game, the opinion of one's partner becomes a critical component of strategy.

When examining the effects of negotiation length and individual differences on user behavior, the ultimatum game provides a helpful testing ground—it is both adequately social and repeatable. A previous study in the literature also fits these two requirements [11]. In that study, subjects were recruited from an online service (Amazon's Mechanical Turk), and were presented with a form of Multi-Issue Bargaining task (as in [7]) where players interacted with agents over multiple rounds. Although a repeated negotiation, the task lasted for only 5 rounds.

In this current work, we extend and improve upon this previous work in three ways: a new population, a new task, and longer interaction. First, we wanted to confirm that the prior effects generalized to other subject populations and tasks, especially since much recent work has found important differences between subject populations [1,2,6]. To do so, the study presented here was conducted with a different population, a simpler task, and a longer timescale than previously used in the literature. We will show that the general design lessons from previous work largely hold.

Further, the expanded timescale of this current work allows for a more nuanced understanding of the dynamics of acceptance over time. By doubling the number of negotiation games to 10, we can detect previously unreported effects that take place after additional negotiations. Furthermore, this work shows that the success of agent strategies is highly dependent on individual differences: namely Social Value Orientation (SVO) and Machiavellianism. Future agents that detect and adapt to

individual differences will therefore be more successful than this current generation.

2.3 Social Value Orientation & Machiavellianism

In addition to observations resulting from the ultimatum game, this work examines two measures of individual differences that are specifically relevant to the negotiation domain. Previous research has identified personality variables, such as SVO and Machiavellianism, that can influence public goods decisions [26]. These measures have been shown to affect outcomes in human negotiation, but there is little agreement on how they affect behavior overtime in negotiation, and even less on their impact for mixed human-agent systems [18,21]. As such, examining their effects in repeated-form human-agent negotiation is largely unexplored.

SVO [22] measures general “pro-self” vs “pro-social” tendencies. Earlier studies of people's social value orientations [9] have demonstrated that, when faced with social dilemmas in which actions that are most personally beneficial conflict with actions that are most beneficial to a larger group or community, many people will decide to pursue communal interests at the expense of their own. Thus, concerns with their own individual outcomes do not appear to uniformly eclipse people's considerations of their broader connections to and responsibilities toward others [4].

Studies [12,22] have consistently confirmed that, compared to those who identify themselves as individualistic, pro-social individuals are more likely to: (a) follow norms of social responsibility that dictate cooperatively sacrificing their own potential gains to improve communal interests in both simulated [3] and real-world [25] social dilemmas [1], and (b) follow norms of equality or fairness that dictate actively seeking to equalize their own and others' outcomes, even when they could easily keep more benefits for themselves [19,21,24].

The second measure, Machiavellianism, quantifies strategic, manipulative, and goal-seeking behavior. Like SVO, Machiavellianism is particularly relevant to the negotiation domain as it has also been found to relate to strategies and behaviors employed in negotiation [27]. Specifically, Machiavellians tend to have similarities with high-narcissism and high-psychopathy subjects [15]. A quintessential Machiavellian might be considered overly rational, detached, or cold, and tend not to be influenced by emotional arguments as easily as those scoring lower on the measure [6]. Further, Machiavellians often hedge their answers when asked direct questions and prefer to obfuscate their desires and goals [27].

Because of these traits, Machiavellians can often gain advantage in certain negotiation and game theoretic tasks, as they tend to frustrate opponents' efforts to recognize beneficial offers. These strategies have been studied primarily in short-term interactions [28], however, and may eventually lose their usefulness in repeated negotiation. Therefore, Machiavellianism is a particularly important measure to examine while studying repeated interactions [27]. Negotiators high in Machiavellianism have also been found to use different language than those lower in

Machiavellianism. This language is important when favors are unenforceable, and must be relayed purely through discussion [5].

In summary, SVO is a helpful measure for discriminating pro-social from pro-self tendencies—whether a person thinks of others or only of him or herself. Machiavellianism, by contrast, is a way of distinguishing how people think of others—in a manipulative and exploitative way, or in a more altruistic one.

3 EXPERIMENTAL DESIGN

3.1 Game Structure

Participants in this experiment were tasked with playing 10 ultimatum games in sequence against a virtual agent (they were told they would be playing against “a computer opponent”). Each ultimatum game was played on an online interface in which a series of cards depicting a specific type of fruit were shown on screen (Figure 1). The experimental goal was to test the effects and antecedents of favor-exchanging behavior. Therefore, each game either featured cards that were vastly more valuable to the player than to the agent, cards that were vastly more valuable to the agent than to the player, or of equivalent value to both parties. In this way, players could win the most points by agreeing to take few cards that were less valuable to them in exchange for the understanding that they would receive more cards that were more valuable to them during the subsequent games.

Specifically, in each game, 20 cards were shown to the player, depicting either an apple, or an orange, or a banana. All 20 of the cards shown in a given game always depicted the same fruit. Players were told before beginning that each orange they received would be worth 2 points to them, while apples and bananas would be worth only 1 point. They were also told that apples were worth 2 points to their opponent, while oranges and bananas would be worth only 1 point to their opponent.

The experiment then followed the game structure as outlined in previous work [11], in which 4 player-receiver ultimatum games are followed by 1 “reversal round” where player acts as the proposer of an ultimatum rather than the receiver. These games are summarized in Table 1. The final 5 games were an exact repeat of this pattern. Thus, the basic structure of [11] was preserved, but the duration of the task was doubled (from 5 games to 10)—and, of course, the task and experimental populations are different.

In the first game, the player was presented with an ultimatum involving apples. Since apples were worth more to the agent than to the player, it was possible for the agent to score up to twice as many points as the player. Game 2 was an orange game, in which the player had the opportunity to score more than the agent. Finally, in Game 5, the player was presented with a banana game, in which the cards were of equal value to both the player and the agent (1 point each). Games 1, 3, 6, and 8 (the “apple” games) were designated “favor-seeking games”, since the agent would be best served by attempting to get most of the items and make it up to the player later. Games 2, 4, 7, and 9 (the “orange” games) were designated “return-opportunity games”, as these games provided the best opportunity for the agent to provide points to the player without giving up much value.

In each game, players were either a receiver or a proposer of an ultimatum deal. When the players were receiving, their agent opponent would select some number of cards to offer to the player (either 5 or 15 cards, depending on the agent and game). As in any ultimatum game, the players were then presented with the option to accept the offer (receiving the points indicated by the offer), or to reject the offer (with both sides receiving 0 points). Along with the offer, they also received a message from the agent, written in text on the screen. There was a still image of the agent’s avatar shown on the screen at all times. When the players were proposing offers, they were given the ability to designate any number of cards to offer to the agent. Once they locked in their proposal, the agent decided whether or not to accept the ultimatum, and points were awarded accordingly. All agents accepted any proposal in which they received at least 25% of the total value (in this case, 5 out of the 20 cards).

3.2 Agent Design and Experimental Manipulation

Agents differed according to two experimental manipulations. Agents were classified as either utilizing favor or non-favor language, and as competitive or cooperative. This 2x2 design allowed the designation of 4 agents: the favor-language cooperative agent (“favor-returning”), the favor-language competitive agent (“betraying”), the non-favor-language cooperative agent (“altruistic”), and the non-favor-language competitive agent (“tough”).

The cooperative/competitive manipulation affected agent behavior, whereas the favor/non-favor manipulation affected only the messages that the agent sent. The favor-language agents (Favor-returning and Betraying) requested favors in the favor-seeking games with language:

“This goal is important to me. I hope you can accept this deal as a favor to me. I’ll really owe you one.”

and on the return-opportunity games:

“Hey, thanks for doing me that favor before. Let me help you out in return.”

Non-favor-language agents (Altruistic and Tough) used more generic language, such as

“I think this deal is acceptable.”

This language is taken from [11].

The behavioral manipulation produced agent behavior that was dependent on the game. Competitive agents (Betraying and Tough), always offered only 5 out of 20 cards, regardless of the game. Cooperative agents (Favor-Returning and Altruistic), however, offered 5 out of 20 cards on the favor-seeking games, but offered 15 out of 20 cards on the return-opportunity games. In this way, they were more “fair”, alternating between claiming the majority of the cards when they were of particular value to them and offering a majority of the cards when they were of particular value to the player. The agents thus vary along two axes—one according to “word” and the other according to “deed”. By exposing participants to one of these four agents, this design allows us to determine both the dynamics of user acceptance as it evolves over all 10 games, and the individual differences of the players that may become relevant.

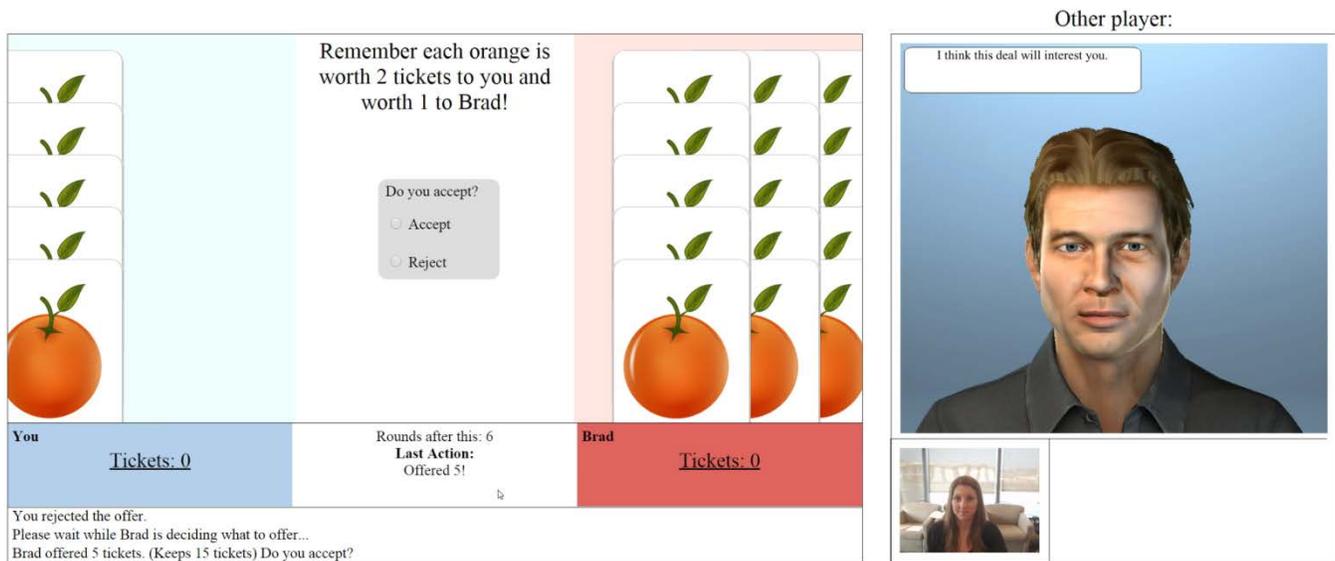


Figure 1: Game interface as seen by the player

Table 1: Game Structure (First 5 games)

Game #	Game type	Offer Made (Cooperative condition)	Offer Made (Competitive condition)	Integrative Potential?
Game 1 (apple)	Favor-seeking	Bad	Bad	Yes
Game 2 (orange)	Return-opportunity	Good	Bad	
Game 3 (apple)	Favor-seeking	Bad	Bad	Yes
Game 4 (orange)	Return-opportunity	Good	Bad	
Game 5 (banana)	User-proposed offer	User choice	User choice	No

3.3 Recruitment and Methodology

The experiment randomly matched participants against one of four agents. 106 participants (53 male, 47 female, 6 did not disclose) were recruited via advertisements on the Craigslist website. They were offered \$30 to spend approximately one hour at our facility, during which they engaged in two studies (one unrelated to the results of this paper). This experiment took approximately 15 minutes, during which the participants were provided instruction in the game, then played through all ten games. Participants were told that they would be awarded a small bonus based on their performance in the game, with higher points scored during the experiment translating into greater chances for rewards. In addition to collecting information about the participant's performance in the game, demographic data was also collected for each participant via a short survey before the games began. This included standard data such as that contained in the US Census (race, age, gender, etc.). During this pre-survey phase, participants were also asked to fill out standard, previously validated individual difference scales designed to determine their SVO and Machiavellianism scores. Following the pre-survey, the participants engaged in all ten ultimatum games in sequence. In

our results, we discuss both the percentage of players that accepted ultimatums in a given game, as well as the total number of cards kept by the player during games in which they were a proposer.

4 RESULTS

When analyzing the results, it is important to note that the quality of offers varied from game to game, as well as by condition. People are more likely to accept oranges than apples simply because oranges are more valuable to them, and acceptance rates are higher in the return-opportunity rounds.

This is true even in the competitive condition, where offers are bad (see Table 1). This baseline difference doesn't matter to the analyses of this section, because the analyses always look at differences between conditions in the same game or game type (favor-seeking vs. return opportunity), and do not compare results from different game types.

We analyze these results in terms of acceptance rates of agent offers. Higher rates mean that the offer was accepted by more participants, and is thus construed as more effective.

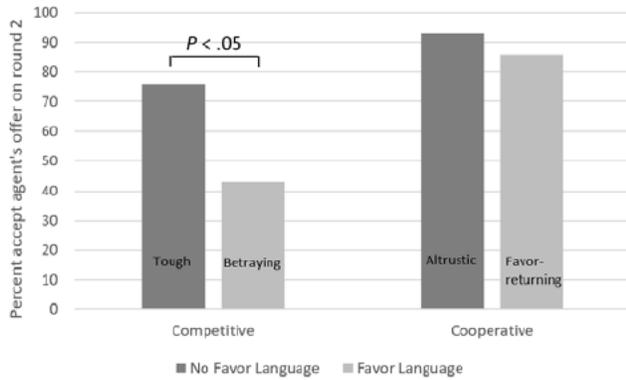


Figure 2: The effect of offer and favor on likelihood of accepting the agents offer in game 2

Table 2: The effect of offer on ultimatum games where agent makes an offer (games 1-4 and 6-9)

Game	Acceptance Rate		Z-value
	Competitive	Cooperative	
1	43%	34%	0.97
2	58%	92%	3.42***
3	17%	40%	2.51**
4	45%	91%	4.33***
6	19%	38%	2.16*
7	47%	94%	4.48***
8	19%	58%	3.95***
9	49%	85%	3.59***

*** is $p < .001$, ** is $p < .01$, * is $p < .05$.

4.1 Agent Performance over Time

To begin, we tested the effect of offer and favor language individually in each game in which the player acted as receiver. For these games, we conducted 2 (offer: competitive or cooperative) \times 2 (favor: favor language or no favor language) \times 2 (accept: yes or no) log-linear analyses. Next, for the reversal games, we conducted 2 (offer: competitive or cooperative) \times 2 (favor: favor language or no favor language) ANOVAs on the amount user kept for themselves.

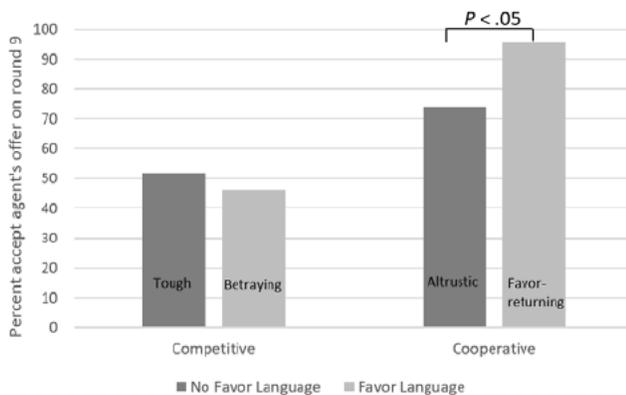


Figure 3: The effect of offer and favor on likelihood of accepting the agent's offer in game 9

Game 1 presented no opportunities for betrayal or returning of favors since it was the first game.¹ By game 2, however, the first game in which favor-language agents either come through with their promise to return the favor (favor-returning) or betray that promise (betraying), we found a backfire effect of betrayal that is similar to the one observed in [11]. Although the interaction and effect of favor could only be considered trend-like at best ($\chi^2 = 1.44$, $p = .23$ and $Z = 1.28$, $p = .20$), with the significant effect of offer ($Z = 3.42$, $p = .001$), the pattern suggests that betraying agents' offers are least likely to be accepted (see Figure 2). Because game 2 was a return-opportunity game where cooperative agents made good offers, agents' offers were more likely to be accepted than rejected ($Z = 4.90$, $p < .001$).

In games 3 through 8, the only significant effect of condition that emerged was the effect of offer (see Table 2). Across all offer games, cooperative agents gain acceptance more than competitive agents ($Z_s > 2.16$, $p_s < .03$). The effect of being cooperative pervaded onto favor-seeking games when these agents made bad offers. However, the likelihood of accepting overall depended on game: in favor-seeking games, agents' offers were more likely to be rejected than accepted ($Z_s > 2.45$, $p_s < .01$), whereas in return-opportunity games (where cooperative agents made good offers), agents' offers were more likely to be accepted than rejected ($Z_s > 3.63$, $p_s < .001$). This oscillating acceptance rate occurred even for competitive agents, again due to the aforementioned baseline difference.

At the end of the interaction in games 9 and 10, the benefit emerged for the favor-returning agent. In game 9, there was a significant interaction between offer, favor, and accept ($\chi^2 = 4.38$, $p = .04$) such that favor-returning agents' offers were most likely to be accepted (as can be seen in Figure 3). This interaction qualified a significant effect of offer on acceptance ($Z = 3.59$, $p < .001$) and a trend for an effect of favor on acceptance ($Z = 1.48$, $p = .14$). As in other return-opportunity games, offers were more likely to be accepted than rejected ($Z = 3.48$, $p = .001$).

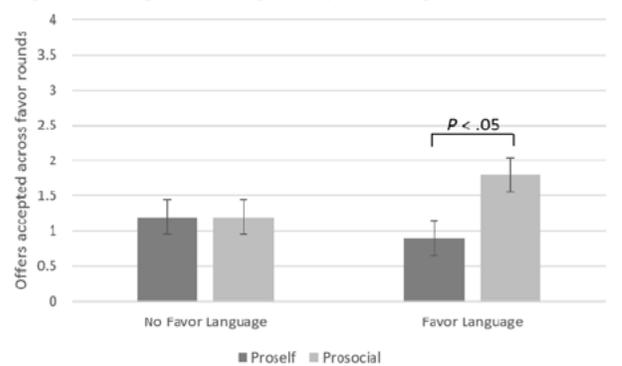


Figure 4: The effect of SVO and favor on acceptance across favor-seeking games

¹In game 1, we tested for failure of random assignment. There was a significant interaction between offer, favor, and accept ($\chi^2 = 9.51$, $p = .002$) such that participants were much more willing to accept offers from the betraying agent than the favor-returning agent. Fortunately, this random variation did not have downstream consequences and does not moderate any of the effects in the subsequent game ($Z_s < 0.99$, $p_s > .32$).

4.2 User's Individual Differences and Agent Performance

To explore which agents performed best based on users' individual differences, we first conducted a 2 (offer: competitive or cooperative) \times 2 (favor: favor language or no favor language) \times 2 (SVO: pro-social or pro-self) ANOVA on the number of times users accepted the agent's offer in favor-seeking games. A marginally significant effect of SVO emerged such that pro-self individuals were less likely to accept the agent's offer in these games ($M = 1.14$, $SD = 1.26$) than pro-socials ($M = 1.51$, $SD = 1.24$; $F(1,84) = 3.00$, $p = .08$); however, this was marginally qualified by an interaction with favor condition ($F(1,84) = 3.64$, $p = .06$). As depicted in Figure 4, favor-language agents perform better with pro-social individuals than pro-self individuals. As above, cooperative agents ($M = 1.71$, $SD = 1.37$) perform better in these favor-seeking games than competitive agents ($M = 0.89$, $SD = 0.97$; $F(1,84) = 8.39$, $p = .005$).

We examined whether these effects held for return-opportunity games. A marginal effect of SVO emerged, but in contrast to favor-seeking games, pro-self individuals were more likely to accept the agent's offer in these games ($M = 3.08$, $SD = 1.34$) than pro-social individuals ($M = 2.70$, $SD = 1.50$; $F(1,84) = 3.65$, $p = .06$). Because cooperative agents made better offers in these games, they did much better than competitive agents ($M = 2.07$, $SD = 1.42$ vs. $M = 3.67$, $SD = 0.91$; $F(1,84) = 37.38$, $p < .001$).

In addition to SVO, we also considered which agents performed best based on players' Machiavellianism scores. We thus conducted a 2 (offer: competitive or cooperative) \times 2 (favor: favor language or no favor language) \times 2 (Machiavellianism: high or low) ANOVA on performance across the favor games. As reported above, we observed effects of offer condition ($F(1,91) = 10.42$, $p = .002$) and trends for interactions between offer and favor ($F(1,91) = 2.47$, $p = .12$). However, effects of and interactions with Machiavellianism could only be considered trend-like at best ($F_s < 1.90$, $p_s > .17$). For the return-opportunity games, however, we found both an effect of offer condition ($F(1,91) = 45.63$, $p < .001$) and a significant interaction with varying Machiavellianism scores ($F(1,91) = 7.03$, $p = .009$). Competitive agents perform worse when playing with users high in Machiavellianism (as depicted in Figure 4). Although the three-way interaction did not approach significance ($F(1,91) = 0.26$, $p = .61$), inspection of the means reveals that this is driven more by users higher in Machiavellianism punishing the betraying agents that failed to return favors in these games. In an analysis across all games, there was also an effect of offer condition ($F(1,91) = 45.72$, $p < .001$), but not of favor condition ($F_s < 1.65$, $p_s > .20$).

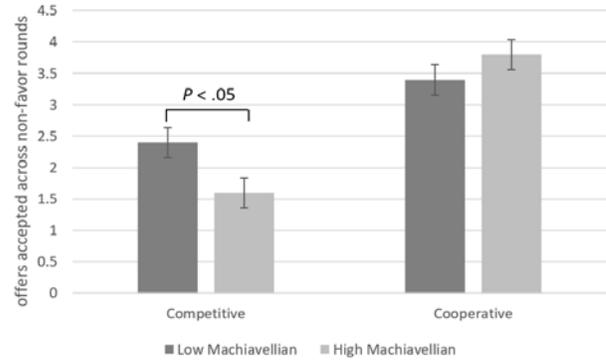


Figure 5: The effect of Machiavellianism and offer on acceptance across return-opportunity games.

5 DISCUSSION

Our work has implications for both general design of agents who engage in repeated negotiations with users, and—importantly—for such agents that are specifically designed for users with certain personalities.

5.1 Designing Agents for Repeated Negotiations

First, the results show that there is a cost of betrayal associated with aggressive early strategies (as shown in literature). Already by game 2, users were reluctant to accept an offer from the betraying agent compared to the agent that was merely competitive (Figure 1). In terms of agent design, agents should not both present poor offers and promise good ones—such a deception will be quickly identified. Rather, performing positively by actually delivering good offers seemed to garner the desired high acceptance rates from users. There did not, however, appear to be any benefit to additionally using favor language (as the favor-returning agent did) in those early games. In short: Figure 2 shows no short-term advantage of using favor-language.

However, the later games paint a far-different story. While users were quick to be cued to the less-than-fair behavior of agents when those agents used favor language, they also eventually detected poor offers regardless of favor language. Furthermore, as Figure 3 shows, by game 9, players rewarded a history of combined good behavior and favor-language by accepting offers from the favor-seeking agent significantly more than other agents. This benefit of cooperation takes some time to create (in this case, 9 games).

These general results themselves extend previous work. We replicate several effects found in literature, and do so using a population with different demographics (Craigslist instead of Mechanical Turk) and different task (Ultimatum instead of Multi-Issue). Moreover, the benefit of cooperation as demonstrated in game 9 is a new effect; by extending the interaction, this work uncovered new temporal effects. There appears to be a tradeoff between short-term and long-term performance that this work begins to quantify. Disreputable actions (such as lying or aggressive techniques) may have immediate gains. But these results show not only that users discover these tactics, but also

that positive actions and congruence between word and deed can have positive outcomes in the long-term.

5.2 Designing Agents for Repeated Negotiations

While general advice on avoiding betrayal costs and gleaned cooperative benefits is helpful, it ultimately is restricted by the myriad individual differences found in negotiators. These “best practices” incrementally push the gold standard for negotiating agents, but do not fundamentally change the algorithms behind them. To this end, we examined the effect of two oft-used individual difference measures to determine the effect of agent strategy on user behavior, in the hopes that there would be effects that advanced agents could exploit.

When examining the SVO of participants, it was found that pro-social users were more likely to grant agent favors than pro-self users. In other words, the social language used by the agent was far more likely to sway pro-social participants (see Figure 4). However, attempting to use this favor language with pro-self users was worse, on average, than simply avoiding using the favor language at all. In short, knowing certain traits about the user can lead to substantial benefits, but guessing incorrectly could backfire. Armed with this knowledge, agents can be designed whose behavior is determined (in part) by models that incorporate such user traits. Given the promise of existing methods of automatically detecting user traits [13], this should be compelling evidence that such lines of research will prove rewarding. By using self-reported measures of individual difference (as done here) or ideally by determining SVO while engaging in negotiation, agents should be able to improve their performance.

The final measure examined was Machiavellianism. Similar to SVO, Machiavellianism score had a significant effect on user behavior when confronted with different agents. Here however, it was the quality of offers (not the language) provided by the agents that caused a behavioral split. As illustrated in Figure 5, users high in Machiavellianism punish competitive agents when those agents have an opportunity to provide good offers. Inspection of the 3-way means revealed that this effect is particularly driven by the betraying agent. Although providing bad offers constantly (the competitive conditions) was shown to provide a constant, negative, main effect throughout the study, high-Machiavellian individuals punish the betraying agent with vindictive gusto, leading to far fewer acceptances in return-opportunity rounds. In this, as with SVO, it is important to know one’s enemy: a good model of the opponent can lead to agents avoiding the pitfall of making overly aggressive offers when dealing with a known high-Machiavellian.

These results demonstrate that the careful design of the strategy of negotiating agents is critical to success. Successful agents must weigh carefully what they know about their opponent and what they know about the structure of the negotiation itself (particularly, how long it will last). Provided with these tools, the next generation of negotiating agents can provide a more robust, effective, and realistic experience when interacting with diverse users.

ACKNOWLEDGMENTS

This work was partially supported by a variety of grants and organizations, which have been redacted for blind review.

REFERENCES

- [1] Kathleen M Alto, Keiko M McCullough, and Ronald F Levant Who is on Craigslist? A novel approach to participant recruitment for masculinities scholarship, 2016
- [2] Christopher Antoun, Chan Zhang, Frederick G Conrad, and Michael F Schober Comparisons of online recruitment strategies for convenience samples: Craigslist, Google AdWords, Facebook, & Amazon Mechanical Turk *Field Methods*, 231-246, 2016
- [3] Daniel Balliet, Craig Parks, Jeff Joireman Social value orientation and cooperation in social dilemmas: A meta-analysis *Group Processes & Intergroup Relations*, 12(4), pp 533-547, 2009
- [4] Linnda R Caporael, Robyn M Dawes, John M Orbell, Alphons JC Van de Kragt Selfishness examined: Cooperation in the absence of egoistic incentives *Behavioral & Brain Sciences*, 683-699, 1989
- [5] Andrea Czibor, Orsolya Vincze, and Tamas Bereczkei Feelings and motives underlying Machiavellian behavioural strategies; narrative reports in a social dilemma situation *International Journal of Psychology*, 49(6), pp 519-524, 2014
- [6] B Fehr, D Samsom, Delroy L Paulhus The construct of Machiavellianism Twenty years later, CD Spielberger, JN Butcher (red), *Advances in personality assessment* pp 77-116, 1992
- [7] Ya’akov Gal, Barbara J Grosz, Sarit Kraus, Avi Pfeffer, Stuart Shieber Colored trails: a formalism for investigating decision-making in strategic environments *Proceedings of 2005 IJCAI workshop on reasoning, representation, & learning in games* 25-30, 2005
- [8] Daniel Kahneman, Jack L Knetsch, Richard H Thaler Fairness and the assumptions of economics *Journal of business*, pp S285-S300, 1986
- [9] Kenneth R MacCrimmon, David M Messick A framework for social motives *Systems Research and Behavioral Science*, 86-100, 1976
- [10] Johnathan Mell, Jonathan Gratch Grumpy & Pinocchio: Answering Human-Agent Negotiation Questions through Realistic Agent Design *In Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems* pp 401-409, 2017
- [11] Johnathan Mell, Gale Lucas, Jonathan Gratch An effective conversation tactic for creating value over repeated negotiations *In Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems* pp 1567-1576, 2015
- [12] David M Messick, Charles G McClintock Motivational bases of choice in experimental games *Journal of exp social psychology*, 1968
- [13] Zahra Nazari, Gale Lucas, and Jonathan Gratch Multimodal approach for automatic recognition of machiavellianism *In Affective Computing and Intelligent Interaction (ACII)* pp 215-221 *IEEE*, 2015
- [14] Martin A Nowak, Karen M Page, Karl Sigmund Fairness versus reason in the ultimatum game *Science*, 289(5485), pp 1773-1775, 2000
- [15] Delroy L Paulhus, Kevin M Williams The dark triad of personality: Narcissism, Machiavellianism, and psychopathy *Journal of research in personality*, 36(6), pp 556-563, 2002
- [16] Joel Ross, Lilly Irani, M Silberman, Andrew Zaldivar, and Bill Tomlinson Who are the crowdworkers?: shifting demographics in mechanical turk *In CHI’10 extended abstracts on Human factors in computing systems* (pp 2863-2872) *ACM*, 2010
- [17] Ramzi Suleiman Expectations and fairness in a modified ultimatum game *Journal of Economic Psychology*, 17(5), pp 531-554, 1996
- [18] Leigh Thompson Negotiation behavior and outcomes: Empirical evidence and theoretical issues *Psychological bulletin*, 108(3), p 515, 1990
- [19] Eric Van Dijk, David De Cremer, Michel JJ Handgraaf Social value orientations and the strategic use of fairness in ultimatum bargaining *Journal of Experimental Social Psychology*, 40(6), 697-707, 2004
- [20] Gerben A Van Kleef, Carsten KW De Dreu, Anthony SR Manstead The interpersonal effects of emotions in negotiations: a motivated information processing approach *Journal of personality and social psychology*, 87(4), p 510, 2004
- [21] Paul AM Van Lange The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientation *Journal of personality and social psychology*, 77(2), p 337, 1999
- [22] Paul AM Van Lange, Ellen De Bruin, Wilma Otten, Jeffrey A Joireman Development of pro-social, individualistic, and competitive orientations: theory and preliminary evidence *Journal of personality and social psychology*, 73(4), p 733, 1997
- [23] Paul AM Van Lange, David De Cremer, Eric Van Dijk, Mark Van Vugt Self-interest and Beyond *Social psychology: Handbook of basic principles*, pp 540-561, 2007

- [24] Paul AM Van Lange, Carsten KW De Dreu Social Interaction: Cooperation and competition Introduction to Social Psychology, 3, pp 341-370, 2001
- [25] Mark Van Vugt, Paul AM Van Lange, Ree M Meertens Commuting by car or public transportation? A social dilemma analysis of travel mode judgements European Journal of Social Psych , 1996
- [26] Mark J Weber, Shirli Kopelman, David M Messick A conceptual review of decision making in social dilemmas: Applying a logic of appropriateness Personality & Social Psych Review, 281-307, 2004
- [27] David Sloan Wilson, David Near, Ralph R Miller Machiavellianism: a synthesis of the evolutionary and psychological literatures Psychological bulletin, 119(2), p 285, 1996
- [28] Yuyu Xu, Pedro Sequeira, & Stacy Marsella Towards modeling agent negotiators by analyzing human negotiation behavior In Affective Computing and Intelligent Interaction (ACII), 2017 Seventh International Conference on Affective Computing and Intelligent Agents pp 58-64, 2017