Varied Magnitude Favor Exchange in Human-Agent Negotiation

Johnathan Mell
USC Institute for Creative Technologies
Los Angeles, CA, USA
johnathanmell@me.com

Gale M. Lucas
USC Institute for Creative Technologies
Los Angeles, CA, USA
lucas@ict.usc.edu

Jonathan Gratch
USC Institute for Creative Technologies
Los Angeles, CA, USA
gratch@ict.usc.edu

ABSTRACT

Agents that interact with humans in complex, social tasks need the ability to comprehend as well as employ common social strategies. In negotiation, there is ample evidence of such techniques being used efficaciously in human interchanges. In this work, we demonstrate a new design for socially aware agents that employ one such technique—favor exchange—in order to gain value when playing against humans. In an online study of a robust, simulated social negotiation task, we show that these agents are effective against real human participants. In particular, we show that agents that ask for favors during the course of a repeated set of negotiations are more successful than those that do not. Additionally, previous work has demonstrated that humans can detect when agents betray them by failing to return favors that were previously promised. By contrast, this work indicates that these betrayal techniques may go largely undetected in complex scenarios.

1 Background

1.1 Motivation

Creating virtual agents capable of realistic, human-like negotiation advances various artificial intelligence tasks and is a key challenge problem across multiple disciplines [15]. Since negotiation is a quintessentially social task, virtual agents that are capable of influencing, understanding, and interacting with human partners must follow and understand social strategies. In order to develop effective social agents, we must design and evaluate new strategies through rigorous empirical study.

Much work has been dedicated to understanding and mapping effective strategies within human negotiation. Indeed, there is a robust literature within the psychology and business communities that showcases the various strategies and effects relevant to human negotiation [3-4][7-10][14][28][30], just as there is a substantial body of work among virtual agent researchers in implementing them [6][11][17-20]. Among these are strategies that deal primarily with *repeated* negotiations, where longitudinal effects such as reputation and favor exchange become relevant. Favor exchange relies on a form of reputation calculus (a mental "ledger" of owed favors) and a perception of how likely a given partner is to return favors. These two concepts are thus entwined, and effective long term agents that participate in repeated interactions must understand and adapt to these human concerns [27].

We are interested in furthering the understanding of favor exchange (and relatedly, reputation), in the human-agent negotiation context. Other work has examined effective negotiation strategies in human-agent competitions [23], as well as the use of favors in simple ultimatum games [21]. However, we are aiming to examine favor exchange in a new and more complex task (multi-issue bargaining).

Beyond merely changing the domain in which these issues are examined, there are further important factors to be disentangled in order to create agents that can accurately model (or at least, effectively interact with) humans. In particular, there has been little work that has performed an analysis of the *magnitude* of favors exchanged. While it is generally accepted that most humans can keep a rough "ledger" of favors owed to others and to themselves, it has been difficult to capture this mental process empirically. If we can measure the sensitivity of humans to various favor exchange paradigms, we can better design agents that can navigate the edges of this strategic technique. Finally, this work serves as an illustrative example of the importance of longitudinal, highly interactive human-agent negotiation. We examine and affect the behavior of human users in human-agent negotiation in a repeated interaction with a virtual agent. By showing that human behavior differs at later points when agent behavior is identical, after humans are exposed at earlier points to divergent agent behavior, we unequivocally state the non-Markovian nature of the human/social-agent interaction problem.

In short, this paper and its contained empirical study answer three separate, but important research items. The study is therefore designed to address three research hypotheses:

a) **favors are an effective tactic** (i.e., they increase the total value claimed when used)

b) **the "magnitude" of the favor returned is/is not important** (i.e., does returning a large amount of value vs. a small amount matter, or is the act of returning any value at all sufficient?)

IVA '20, October 19–23, 2020, Virtual Event, Scotland Uk

ACM ISBN 978-1-4503-7586-3/20/09…$15.00
https://doi.org/10.1145/3383652.3423866*https://doi.org/10.1145/3383652.3423866*

c) **negotiation history will affect human behavior even if current agent behavior is identical** (i.e., two agents who act identically in the current interaction may still yield different results due to their differing histories)

We specifically define "negotiation" in this work as being akin to the "multi-issue bargaining task" that is common in the literature (e.g., as described in [12]). In this task, two parties attempt to divide up a finite number of issues, allocating some number of these issues to either side. Each issue may have quantities greater than one, allowing each issue to be split partially between each side. Moreover, each issue is worth an unknown amount to each side, and these amounts may be discovered during the course of the negotiation. Both parties also have a "best alternative to negotiated agreement" (BATNA), which represents their score (normally a small number) if they fail to reach agreement with their partner. The "negotiation space" refers to the entire discrete space of fully-allocated solutions, wherein all items have been distributed to one side or the other, and none remain unallocated.

1.2 Trust and Favors

To study favor exchange, we must contend with other critical negotiating components—reputation and trust. We make the argument that the likelihood of returning favors—and thus, their expected value—is based on the level of trust in the opponent. Trust can be developed very swiftly in social situations [25], but is of course also affected by the generalized reputation an opposing negotiator may have. Reputation may be seen as a public concept (e.g., Yelp review ratings) or as an individual record (e.g., personal notes/memories about an interaction). We focus on the latter sense in this work, and how the exchange of favors can alter trust (and individualized reputation) in meaningful ways.

Indeed, not only does information about reputation potentially provide information about how an opponent may conduct themselves before negotiations begin, savvy negotiators can manipulate their own reputation in the same way to achieve an advantage. For example, if a negotiator is seen as trustworthy before a negotiation begins, then the opponent is less likely to mistrust potentially valuable information provided about preferences. A negotiator who is preparing to face off against an opponent who is known to be very tough may be more guarded, but may also come in ready to concede in order to preempt a long and vicious fight.

But, reputation is more than just an *a priori* baseline on which to base initial strategies in the absence of real information. Since reputation is assumed to be dynamic (what you do will change your reputation), negotiators must select strategies that will not only lead them to success in a single negotiation, but will also lead to the desired reputation. This assumes that there is a chance of interacting with the same individual again (otherwise, reputation gains would not help to gain future value).

As a case study in this phenomenon, we examine work by de Melo et al [10]. In this study, human participants were instructed to provide information about how they would negotiate in the future, by providing instructions to an agent that served as their representative. Participants who acted through a representative—as compared to people who acted directly with no intermediary—construed the problem on a higher (relationship and norms) level, and selected fairer behavior for their agent. This indicates an increased concern with concepts of reputation.

Still, the dynamics of trust and reputation are not fully understood. In other work [22], it has been shown that while this initial consideration for fairness may exist, prior negative negotiation experience leads to increasingly manipulative behavior, as participants begin endorsing techniques like lying and negative emotion use. However, since this previous study does not make explicit the expectation of repeated future interactions with the same partner, it is unclear how relationships between the human and various agents would develop over time.

In short, humans manage to maintain a concept of relationships, especially with regards to tasks with clear outcomes. Negotiation is one such clear-outcome task. In this work, we examine a narrow strategy related to reputation that is effective in human negotiation called "favors and ledgers". This strategy allows negotiators to accept unfair outcomes in the short term with the expectation that these favors will be repaid over time, thus unlocking greater shared value ("growing the pie"). Related work has shown this strategy to be effective in human-agent interactions, but in limited contexts (not full-fledged multi-issue negotiation) [21]. Similar work has been conducted, focusing on trust [13] and social dependencies [16].

1.3 Pareto-Optimality over Time

One of the key areas for expanding research within human-agent negotiation is the study of temporally-aware agents. Temporally-aware agents have benefits that are most clearly illustrated by examining the additional value that can be claimed across multiple integrative negotiations. These benefits have been shown in ANAC 2018 [23], and provide the mathematical motivation and explanation for why favor-exchanging strategies are effective and sought-after.

Within a given negotiation, division of resources between competing sides can be represented graphically by the set of points representing the utility that each participant receives from a given distribution. Each point that does not generate strictly less utility for both parties is considered to be Pareto optimal (lying on the Pareto frontier). Formally, given a set S of points representing the joint utility of a deal, the set of Pareto optimal points P is defined as:

$$P = \{p\} \mid \forall\, p \in S, \nexists\, q \in S, (p_x < q_x \wedge p_y < q_y)$$

Thus, points falling below the curve generated by these points are considered suboptimal (or "inefficient"), as the deal that those points represent could be improved for one player without harming the other.

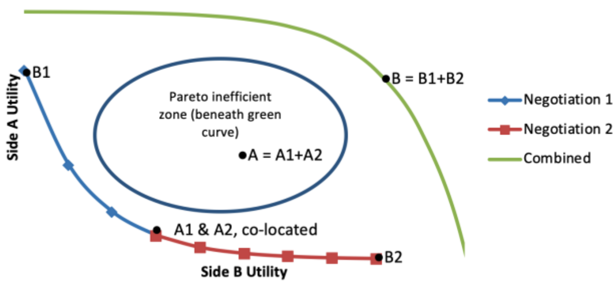

Figure 1: The space of negotiation options in a simple, two-negotiation repeated interaction

Unfortunately, when repeated negotiations are allowed to occur, simply combining Pareto optimal solutions in each individual negotiation can be arbitrarily inefficient over time. This is clearest when the Pareto frontier is convex (Figure 1). In this figure, two negotiations are represented by the red and blue curves. The green curve represents the most extreme edge of the sum of the x- and y-values of these individual negotiations, and thus the area below that curve represents the total problem space of all possible deals in the red and blue negotiations, taken consecutively. In short, this green curve represents the optimal total value of the outcome of two consecutive negotiations.

In this example, humans are likely to pursue the "fair solution" in either individual negotiation (an even split, illustrated as the deal solution at points "A1" and "A2" in Figure 1, which exists on both the red and blue curves). This obvious solution may be efficient for the individual negotiation, but if this even-split solution is chosen for both the red and blue negotiations, it will lead to a summed solution for the combined negotiation that is well below the Pareto optimal zone. In other words, two even splits in a row will sum to fall in the "Pareto inefficient" zone, shy of the green curve, at point "A". Conversely, choices "B1" and "B2" are possible deals in the red and blue negotiations (respectively), and they are *also* efficient when summed (i.e., their sum lies on the green Pareto frontier, at point "B"). However, humans are unlikely to choose either of these "B" points (without prior training) as they would be seen to be violating the norm of fairness for an individual game. Thus the choice of either "B" point is somewhat counterintuitive, but in fact the two combine to form a Pareto efficient *and fair* solution over two consecutive games. Formally, this two-negotiation solution has been defined as [21]:

$$P2 = \{p1 + p2\} \mid \forall\, p1, \in S1, \forall\, p2 \in S2, \nexists\, q1 \in S1, \nexists\, q2 \in S2$$

$$(p1_x + p2_x < q1_x + q2_x \wedge p1_y + p2_y < q1_y + q2_y)$$

In terms of effective—and simplified— negotiation advice, this example shows that choosing "unfair" solutions in single negotiations can lead to "growing the pie" while remaining fair in multiple, repeated negotiations.

Table 1: Example Integrative Issue Utilities

	Apples	Bananas
Item Quantity	4	4
Item Utility to A	3	1
Item Utility to B	1	3

Favors and ledgers is one approach of social interaction that allows parties to discover and achieve such efficient (yet unintuitive) solutions, by recognizing the implications of the changing utilities from negotiation 1 to negotiation 2.

Next, consider for example the simple utility table shown in Table 1. In this example, it is a common but nonoptimal solution to evenly split all items between the two sides. In actuality, however, both sides could do better by simply giving Side A all of the apples, and Side B all of the bananas. If one side recognizes that there will be surplus of apples in the future, that side may agree to forego a "fair" split of the bananas today, with the promise of receiving a similar favor in some future negotiation. This technique of "banking" joint value can be very effective in negotiations, and helps establish mutually beneficial relationships between negotiation partners. Even in situations where payoffs are uncertain (one side does not know the other prefers bananas), exchanging favors is still a viable strategy. Of course, malicious manipulation of favor returns, in which one party claims to incur a favor by accepting a poor deal when in fact it was a good deal for them, is also possible.

In the case of repeated negotiations, achieving Pareto-optimal-over-time results can be accomplished more easily through the repeated exchange of favors that result in *locally* unfair results but *overall* fair exchange. We can therefore predict that virtual agents that can maintain an internal state of their own "ledger" while taking into account opponents' "ledgers" will allow them to robustly deal with these temporal considerations.

The study presented in this paper will examine the use of favors and ledgers in a human-agent negotiation. Specifically, it will examine if the tactic is useful in gaining points over time in practice, and to what extent humans can reason over favor magnitude and favor history, in a fidelitous, dynamic interaction.

2 Experimental Design

2.1 Study Design

Since the nature of favor exchange relies on multiple interactions over time, we design a repeated negotiation scenario, where humans and agents will engage in several back-to-back negotiation rounds. We also customize the favor-returning behavior of a virtual agent. We choose to design both the chosen negotiation tasks as well as the agents themselves using the IAGO negotiation platform, a web-based API designed to create and measure human-agent negotiation interactions [20]. IAGO provides several channels for interacting with a virtual agent—namely, a virtual negotiation table with movable objects (Figure 2, left middle), an embodied agent that can display emotion (Figure 2, left top), and a text-based chat-history where players can send pre-scripted responses (Figure 2, right).

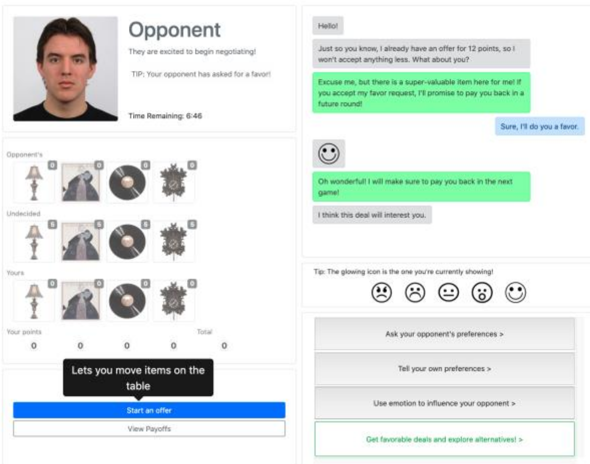

Figure 2: IAGO Research Platform

In order to examine our research questions relating to the magnitude of favors, our primary independent variable is the size of returned favors. We create a 3-cell, between-subjects experimental design, with an additional non-matched 4th cell. The study involves an IAGO agent with three variants of favor-exchanging behavior: the agent will either ask and return favors (favor-reciprocating), ask for favors but weakly return them (demanding), or ask but not return favors (betraying). The final agent makes no explicit calls for favors (no promise), and is therefore included to evaluate the efficacy of favor behavior at all (hypothesis “a”). Agents try to secure the most points in 3 back-to-back 7-minute negotiations.

There are structural differences that promote the effective use of favor exchange for both sides. The first negotiation features a structure that has high-value items for the agent, but low-value items for the human. The second negotiation contains a reversed structure (high-value for human, low-value for agent). The final negotiation is structurally equal for both sides, with the items generally worth few points. This structure is similar to that of other work [23], and also provides an incentive to readily accept the favor in round 1 because there is an obvious structural basis to do so.

All agents pursue an aggressive strategy in round 1, but the favor-reciprocating, demanding, and betraying agents justify their behavior by claiming that a favor will be paid back later. The no-promise agent does not make any claims as to its future behavior and does not use favors. In round 2, the agent negotiates aggressively if it is a betraying or no-promise agent, negotiates on nearly-fair terms if it is a demanding agent and a favor is owed from round 1, and gives ground if it is a favor-seeking agent and a favor is owed from round 1. This varied behavior allows us to directly evaluate hypothesis “b”. In the final negotiation, all agents pursue a fair, consensus-building strategy. In this way, we are able to examine the results of the history of their behavior leading up to round 3, while their behavior within round 3 remains identical (thereby allowing us to evaluate hypothesis “c”).

This study therefore features agents that explicitly attempt betrayal as well as those that promote sincere descriptions of their future behavior. We can narrow our research goals further:

It is hypothesized, based on prior work in simpler domains [24] that there will be both a benefit of cooperation for the favor-seeking agents in the third negotiation, as well as a cost of betrayal for the betraying agents. However, demanding agents may or may not be viewed as returning the favor adequately in round 2, and therefore their performance is as of yet unclear. It is further hypothesized that the three agents that request favors in round 1 will result in higher acceptance rates than the agent that does not (the no-favor agent). The results of this study will confirm which strategies, in general, are the most effective ones for agents to employ when negotiating over an extended period of time. The behavior of the agents is summarized in Table 2.

This study was conducted on Amazon’s Mechanical Turk (MTurk) service, with N=161 subjects recruited. Best practices were followed, including tutorials, attention check questions, recruitment criteria (e.g., high worker rating), and allotment of lottery entry for a cash payment for high performance across the negotiation. After filtering for attention check failure and user absence (repeated timeouts), we retained N=105 (“Gothel” N=23, “Jiminy” N=23, “Gaston” N=20, “Ursula” N=39). All study procedures were approved by a university review board for ethics.

2.2 Agent Design

We are able to design favor-exchanging virtual agents through the IAGO platform, which provides an API for virtual negotiating agent design [20]. Indeed, IAGO facilitates the development of these kinds of agents by providing session-long user states that allow agents to be designed that can recall information from previous interactions. All agents used in this study utilized the same kinds of emotional behavior (responding with positive emotion to good events, and never using anger). They also used the same embodiment of the agent: a photorealistic male. We note that varying the gender of the agent may cause differences in overall rapport (e.g., see [19]), so we kept this constant in the design.

Within the IAGO platform, favor exchange is explicitly supported through a subtype of message events. Both human players and agents are capable of expressing the key events required to update and maintain a ledger of favors. Specifically, both sides can request favors, accept and reject these requests, and explicitly (claim) to return favors. All of these are non-binding convenience communications—humans and agents both maintain their actual ledgers internally. In particular, this means that claiming to return a favor is untied to the truth of the actual return (something often better expressed through the “deeds” of the actual offers received).

Table 2: Agent Behavior

	Round 1	Round 2	Round 3
Structure supports?	Agent side	Human side	Neither side
Favor-reciprocating agents ("Jiminy")	Favor request	Return large favor opportunity	No favor request; favor grant possible
Betraying agents ("Gothel")	Favor request	No favor return	No favor request or grant
Demanding agents ("Ursula")	Favor request	Return small favor opportunity	No favor request; favor grant possible
No-promise agents ("Gaston")	No favor request	No favor return	No favor request or grant

Implementation of the favor behavior includes a new set of dialog options within IAGO to discuss favor requests and returns. In particular, all favor-utilizing agents always open round 1 with a favor request. If that request is accepted verbally, it leads to an actual favorable offer. If the offer is then accepted, then agent's ledger is updated accordingly. Betraying agents ignore their own ledger (so they never attempt to return favors). Favor-reciprocating agents and demanding agents both try to pay back any incurred favors but do so in different magnitudes (the reciprocating agents offer an entire issue of items for free, while the demanding agents offer a single item for free). Both reciprocating agents and demanding agents will grant users favors if asked, but only if they are not already owed a favor (and not in round 1, where the structure favors the agents).

While we believe this design yields results that improve our understanding of human-like negotiating agents, the design does have some limitations. The agents try to take into account any preferences that users may state, and also must assess whether or not the users actually respond positively to the favor requests in round 1. Since interactive agents must be consistent for results to be meaningful, this means that the experience within an experimental cell may differ slightly from subject to subject. For example, it is problematic for a favor-reciprocating agent to return a favor in round 2 if it did not grant one in round 1. Similarly, the user experience may differ slightly in round 3, depending on whether or not users ask for favors from the agent. Still, we argue that these divergences from strict reproducibility are necessary in order to provide an ecologically-valid and truly interactive experience.

The designed agents are therefore capable of altering their strategy to the individual situation, as both favor-reciprocating and demanding agents will only return favors if they owe the human a favor, OR if directly asked when their ledger is neutral. Although this essentially complicates the analysis for agents that were engaging in favors vs. those that did not, we collapse across this distinction in the pursuit of realistic, ecologically-valid, and interactive behavior. Agents are fundamentally dynamic and adaptive in their design. For example, all agents have an internalized conception of "fairness" based on their internal mental model of the opponent. These agents look for a moderately sized positive margin over their opponents in all deals. However, adverse events (such as offer rejections) will reduce this margin over time.

3 Results & Discussion

We tested for differences between agents in their negotiation outcome. We conducted a one-way ANOVA on total points received by the agents across all three rounds. In the omnibus ANOVA, the agents significantly differed on the number of points they earned across the three negotiations ($F(1, 97) = 2.77$, $p = .045$). These results are summarized in Figure 3, and indicate that there was an overall significant effect of favor behavior.

We performed planned post-hoc analysis to determine the agents driving this effect. Follow-up planned contrasts revealed that, while the agents that used favor language (Gothel, Ursula, and Jiminy) on average received more points than the agent that did not use favor language (Gaston; $t(97) = 2.58$, $p = .011$); the agents that used favor language did not significantly differ among each other ($ts < 1.16$, $ps > .25$).

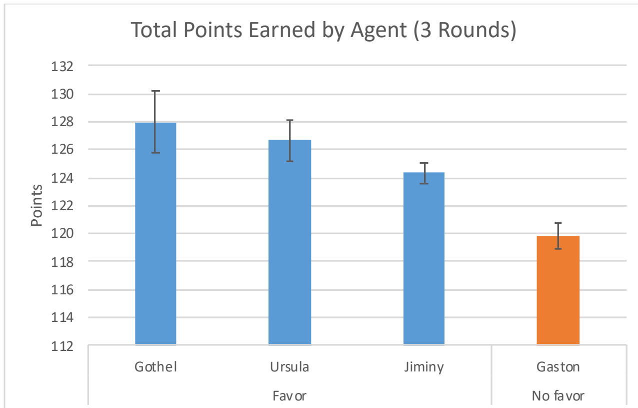

Figure 3: Points Earned in Total (favor agents in blue). Planned 3v1 (blue v orange) contrasts are significant; no other significant differences.

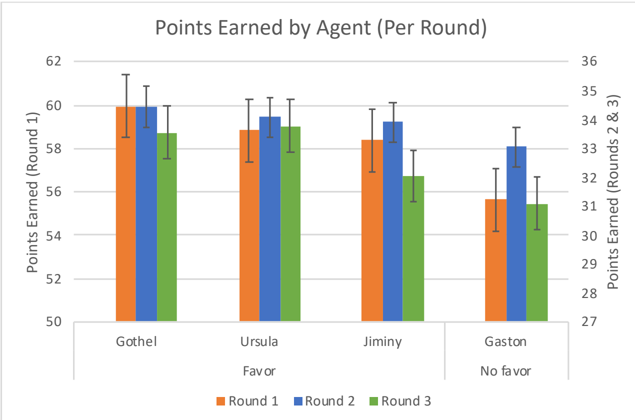

Figure 4: Points Earned Per Round

Breaking this down by negotiation (Figure 4), we find that the primary difference between favor-utilizing and non-favor-utilizing agents is driven by the results in round 1. Analysis by round revealed that this overall difference tended to be driven—as expected—by round 1 ($F(1, 97) = 1.60$, $p = .19$). Due to the behaviors of the agents in asking for favors as well as the structure of the round 1 negotiation favoring the agent, this result clearly indicates that merely asking for favors is an effective technique, which supports our first research hypothesis "a".

We also found that, while not significant, round 3 ($F(1, 97) = 2.35$, $p = .077$) also contributes to this difference somewhat, but there was no hint of differences between agents on round 2 ($F(1, 97) = 0.72$, $p = .55$). Even though the omnibus effect did not reach significance for round 3, follow-up planned contrasts showed the same pattern; while the agents that used favor language on average received more points than the agent that did not use favor language ($t(97) = 1.96$, $p = .05$), the agents that used favor language did not significantly differ among each other (ts < 1.56, ps > .12). The same pattern also emerged for round 1: again, the agents that used favor language on average received more points than the agent that did not use favor language ($t(97) = 2.06$, $p = .04$), but none of the differences *between* the agents that used favor language approached significance (ts < 0.76, ps > .44). On round 2 though, the comparison of favor to non-favor agents did not reach significance ($t(97) = 1.36$, $p = .18$), and again none of the contrasts between favor agents approached significance (ts < 0.55, ps > .58).

While the differences among the three favor-granting agents do not reach traditional levels of significance, the betraying agent, "Gothel" does trend toward being the highest-scoring agent. In previous work, [24], there has been shown to be clear cost of betrayal that appeared early in a set of negotiations. Here, this difference does not appear, and, if anything, appears to be reversed. One possible explanation for this difference is the relative complexity of the task—an IAGO-driven full multi-issue bargaining task is far removed from repeated ultimatum games found in prior work. Indeed, while the favor results indicate that people are indeed capable of perceiving that favors are being asked, participants may not be able to grasp that they are being outmaneuvered by the betraying agent. Our second research hypothesis is therefore somewhat inconclusive, as the magnitude of the favor returned does not show significant differences.

To examine our final research hypothesis, we turn to analysis of the round 3 results among agents that have near-identical behavior in that round (but differ according to their historical behavior). This is accomplished by comparing the betraying (Gothel) and no-favor agents (Gaston). One-way analysis of these two agents in Negotiation 2 indicates a significant difference ($t = 2.281$, $p = 0.029$, no variance assumption). This result indicates that indeed, negotiation history is critically important in reaching conclusions about socially-aware agent behavior and design.

These results seem to indicate that favor-exchange behavior is certainly perceivable by human users in repeated multi-issue bargaining tasks. The usefulness of favors is demonstrated, although the costs of failure to return is still unclear. Part of this mitigated cost of betrayal may be due to the complexity of the task.

We also note that most of the participants did engage with the favor behavior of the agents that demonstrated it. This included generating favor requests of their own (13-15% of participants per round), as well as simply responding to the agent's favor requests.

In order to further investigate the extent to which participants were aware of their own favor-granting behavior, we examined the results of a self-reported questionnaire that took place between rounds. In this questionnaire, participants were asked whether or not they "did the agent a favor" in round 1.

We performed additional analysis comparing this self-report measure of favor acceptance in round 1 to the actual behavior of the human participants. It was found that respondents to the self-report question believed they were giving a favor in much higher quantities than they actually did so using the interface. Notably, some of these self-reports came from Gaston, the no-favor agent that never explicitly asked for favors! Therefore, the non-favor agents are still *perceived* as being given favors. This result reinforces the idea that human perceptions of favors are highly mutable. Furthermore, these results are based on actual interactive data, and are thus arguably more ecologically-valid than other methods.

4 Conclusions & Future Work

By its nature, work performed at the intersection of human psychology and virtual agent design performs twin goals. Namely, it provides us insight about how humans work in the wild (especially when we use tools that preserve the ecological validity of such interactions, like IAGO). Additionally, it allows us to deploy and test virtual agents that can effectively work within these goals to yield the results we want—whether those results be victory, education of participants, or simply efficient interaction. This study provides valuable benefit to both goals.

Our results show the clear effectiveness of favor exchange in multi-issue bargaining tasks—a new domain for this strategy. Previous work in human-human negotiation has proven the effectiveness of the strategy, but we must rigorously confirm these results in a human-agent context, where agents are often considered differently than their human analogues. For psychologists, this provides clear evidence that social techniques in negotiation, when employed by virtual agents, maintain their vigor. For agent designers, this result shows a roadmap for designing advanced strategies in new agents designed to work with humans.

Beyond this mere demonstration of effectiveness, we have contributed new knowledge that humans may not often engage in the rigor needed to evaluate the magnitude of favor returns. Or, at the very least, they may rely on heuristics and Type 1 thinking [2] without triggers to make situations warrant further examination. For those interested in teaching humans to be better negotiators, this point—that special attention should be paid to favor magnitudes—is therefore of special interest. We also note that while the true value of the favor is known to us in this experiment, the *perception* of the value of the favor is not directly measurable, and may be systematically influenced by the agent behavior. Indeed, this hidden potential mediator makes further analysis of favor magnitude even more difficult—we encourage future work to address this, potentially through participant self-report.

Finally, we have reiterated the point shown by other researchers that repeated interactions present a new class of more complex problems, and that these problems require a lengthy, past-looking analysis of events, rather than a reliance on static state machines. While this non-Markovian notion is not new (nor, would we argue, cause for alarm), it is clearly demonstrated by the difference in participant behavior in round 3 of the study—a time when all agents act similarly but have systematically-differing histories.

Our study has some limitations—as we pursue ever-more interactive behavior on behalf of our agents, we necessarily sully the purity of our empirical design. As every participant has a slightly different experience with our agents, their mental models and interpretations of those agents' behaviors necessarily differ. Regardless, we encourage this type of experimental work, as any differences we do find are therefore more likely to be repeatable and impactful.

In the future, we wish to pinpoint the line between betrayal and cooperation more precisely. What triggers are there that will allow humans to detect and monitor "good" and "bad" favor returns? What lines can an agent cross before its reputation and the trust required for favor exchange is damaged? And, is that damage irreparable? These and other questions tie greatly into research on human machine trust, and we welcome extensions by researchers in those subfields, and others [5][13][26]. For now, we will continue to develop social agents for repeated negotiation, and aim to further the development of such agents that use favor exchange and other hitherto essentially human strategies.

ACKNOWLEDGMENTS

Portions of this work have previously appeared in the primary author's doctoral dissertation [1]. The content has been adapted and new analyses included for the current submission.

This research was sponsored by the Army Research Office and was accomplished under Cooperative Agreement Number W911NF-20-2-0053. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

[1] Mell, J. University of Southern California Ph.D. Dissertation, 2020. A Framework for Research in Human-Agent Negotiation. ProQuest Dissertations Publishing.

[2] Allen, Andrew .P., and Thomas, Kevin E., 2011. A dual process account of creative thinking. *Creativity Research Journal* 23, no. 2, pp.109-118.

[3] Bazerman, M. H., & Neale, M. A. 1993. *Negotiating rationally.* Simon and Schuster.

[4] Blascovich, J., Loomis, J., Beall, A. C., Swinth, K. R., Hoyt, C. L., & Bailenson, J. N., 2002. Immersive virtual environment technology as a methodological tool for social psychology. Psychological Inquiry, 13(2), 103-124.

[5] Castelfranchi, C. and Falcone, R., 1998, July. Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In *Multi Agent Systems, 1998. Proceedings. International Conference on* (pp. 72-79). IEEE.

[6] Crandall, J. W., & Goodrich, M. A., 2005. Learning to teach and follow in repeated games. In *AAAI workshop on Multiagent Learning*.

[7] Curhan, J.R., Elfenbein, H.A. and Xu, H., 2006. What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *Journal of personality and social psychology*, *91*(3), p.493.

[8] De Dreu, C.K., Koole, S.L. and Steinel, W., 2000. Unfixing the fixed pie: a motivated information-processing approach to integrative negotiation. *Journal of personality and social psychology*, *79*(6), p.975.

[9] de Melo, C. M., Carnevale, P. J., Read, S. J., & Gratch, J., 2014. Reading people's minds from emotion expressions in interdependent decision making. Journal of personality and social psychology, 106(1), 73.

[10] de Melo, C. M., Marsella, S., & Gratch, J., 2016, May. Do as I say, not as I do: Challenges in delegating decisions to automated agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*(pp. 949-956). International Foundation for Autonomous Agents and Multiagent Systems.

[11] Faratin, P., Sierra, C., and Jennings, N.R., 1998. Negotiation decision functions for autonomous agents. In *Robotics and Autonomous Systems* 24, no. 3-4: 159-182.

[12] Fatima, S.S., Wooldridge, M. and Jennings, N.R., 2007, May. Approximate and online multi-issue negotiation. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems* (pp. 1-8).

[13] Fulmer, C.A. and Gelfand, M.J., 2013. How do I trust thee? Dynamic trust patterns and their individual and social contextual determinants. In *Models for intercultural collaboration and negotiation* (pp. 97-131). Springer, Dordrecht.

[14] Gillespie, J. J., Thompson, L. L., Loewenstein, J., & Gentner, D. 1999. Lessons from analogical reasoning in the teaching of negotiation. *Negotiation Journal*, *15*(4), 363-371.

[15] Gratch, J., DeVault, D., Lucas, G. M., & Marsella, S. 2015, August. Negotiation as a challenge problem for virtual humans. In *International Conference on Intelligent Virtual Agents* (pp. 201-215). Springer, Cham.

[16] Grosz, B.J., Kraus, S., Talman, S., Stossel, B. and Havlin, M., 2004, July. The influence of social dependencies on decision-making: Initial investigations with a new game. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2* (pp. 782-789). IEEE Computer Society.

[17] Hindriks, K., Jonker, C.M., Kraus, S., Lin, R. and Tykhonov, D., 2009, May. Genius: negotiation environment for heterogeneous agents. In *Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 2* (pp. 1397-1398). International Foundation for Autonomous Agents and Multiagent Systems.

[18] Koley, G. and Rao, S. 2018, Adaptive Human-Agent Multi-Issue Bilateral Negotiation Using the Thomas-Kilmann Conflict Mode Instrument.
[19] Krämer, N.C., Karacora, B., Lucas, G., Dehghani, M., Rüther, G. and Gratch, J., 2016. Closing the gender gap in STEM with friendly male instructors? On the effects of rapport behavior and gender of a virtual agent in an instructional interaction. *Computers & Education*, *99*, pp.1-13.
[20] Mell, J. and Gratch, J., 2017, May. 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). International Foundation for Autonomous Agents and Multiagent Systems.
[21] Mell, J., Lucas, G., Gratch, J., 2015. 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). International Foundation for Autonomous Agents and Multiagent Systems.
[22] Mell, J., Lucas, G., Gratch, J., 2018. Welcome to the Real World: How Agent Strategy Increases Human Willingness to Deceive, In *Proceedings of the 2018 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems.
[23] Mell, J., Gratch, J., Aydogan, R., Baarslag, T., and Jonker, C.M. 2019. "The Likeability-Success Trade Off: Results of the 2nd Annual Human-Agent Automated Negotiating Agents Competition", In *Proceedings of the 8th International Conference on Affective Computing & Intelligent Interaction*.
[24] Mell, J., Lucas, G., Mozgai, S., Boberg, J., Artstein, R. and Gratch, J., 2018, November. Towards a Repeated Negotiating Agent that Treats People Individually: Cooperation, Social Value Orientation, & Machiavellianism. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents* (pp. 125-132). ACM.
[25] Meyerson, D., Weick, K.E. and Kramer, R.M., 1996. Swift trust and temporary groups. *Trust in organizations: Frontiers of theory and research*, *166*, p.195.
[26] Olekalns, M. and Smith, P.L., 2009. Mutually dependent: Power, trust, affect and the use of deception in negotiation. *Journal of Business Ethics*, *85*(3), pp.347-365.
[27] Pruitt, D.G., 2013. *Negotiation behavior*. Academic Press.
[28] Syna Desivilya, H., & Yagil, D. 2005, January. The role of emotions in conflict management: The case of work teams. In *IACM 17th Annual Conference Paper*.
[29] Trope, Y. and Liberman, N., 2010. Construal-level theory of psychological distance. *Psychological review*, *117*(2), p.440.
[30] Tversky, A., & Kahneman, D. 1981. The framing of decisions and the psychology of choice. *science*, *211*(4481), 453-458.