Reducing Mistrust in Agent–Human Negotiations

Yinping Yang, *Institute of High Performance Computing, A*STAR, Singapore*

Horacio Falcão, *INSEAD*

Nuno Delicado, *Pluris Pte. Ltd.*

Andrew Ortony, *Northwestern University*

Can a software agent alleviate a human counterpart's mistrust by tactically disclosing additional information during a negotiation? Experimental results indicate that it can, and that more and better agreements result.

A primary challenge in designing intelligent agents to serve as surrogates for human decision makers in negotiation contexts is that there's always *incomplete information* between an agent and its counterpart,[1] especially in situations that present the potential for tradeoffs among multiple issues. An automated selling agent can know its own preferences relative to the issues of, say, price, quantity, and delivery options, but—as with humans—it can at best only surmise the preferences of its counterpart and, without knowing those preferences, it might fail to recognize tradeoffs that offer a better outcome.

Most existing work on automated negotiation has focused on *agent–agent* negotiations and has addressed the incomplete information problem using a variety of machine-learning techniques to make inferences about counterparts, including their resistance points, preferences, and possible tradeoffs.[2-4] Such work in agent–agent negotiations hasn't needed to deal with how to give negotiation agents the kind of social intelligence that facilitates strategically and socially advantageous interactions in human–human encounters. However, as automated negotiation becomes more widespread, a software agent acting on behalf of its principal might well find itself negotiating with a human rather than with another software agent.

In these cases of *agent–human* negotiation,[5,6] even a modicum of social intelligence is likely to result in the agent negotiating more effectively, enhancing both the economic impact of its activities and the social-psychological and affective consequences for its human counterpart. The work that we describe here represents a first step toward giving an agent in an agent–human negotiation a way to reduce the mistrust that's often found in online negotiations.

Mistrust in Negotiations

One thing we know about humans when they engage in online negotiations is that they tend to trust less in this setting than when they negotiate face to face.[7] Because negotiators' interests are often in conflict, it can be difficult for a well-intentioned negotiator—human

1541-1672/14/$31.00 © 2014 IEEE
Published by the IEEE Computer Society

or otherwise—to send a win–win signal or to interpret such a signal from a counterpart as sincere. Two fundamental dilemmas explain why this is the case[8-9]:

- The *dilemma of honesty* concerns how much true information one party should reveal to the other. Revealing everything about his/her preferences makes a negotiator vulnerable to exploitation by the other party, whereas telling nothing increases the likelihood of ending up at an impasse.
- The *dilemma of trust* relates to the extent to which a negotiator should believe the other party. Believing everything can lead to exploitation, while believing nothing makes it difficult to reach an agreement.

One approach to ameliorating these information disclosure problems is for negotiators to adopt a win–win negotiation strategy by proactively sharing some of their interests.[10] The idea is that, by increasing the amount of shared information, the parties are more likely to identify opportunities for preference tradeoffs, thus resulting in better agreements and higher mutual satisfaction. It's also likely that such gestures, especially if reinforced by subsequent behavior, can reduce any existing mistrust between the negotiating parties.

Another approach that can lead to better outcomes in multi-issue negotiations is for negotiators to use multiple equivalent simultaneous offers, wherein offer alternatives that have equivalent utility for one party can have different utilities for the other party.[11] For example, a vendor might propose to a customer a price of $99 for a product delivered in two weeks, or $119 if delivered in three days. By simultaneously giving such offers (assumed to be of equivalent utility for the vendor), a negotiator signals flexibility while at the same time increasing the chance that a particular offer will appeal to the counterpart. Laboratory studies of human–human negotiations indicate that multiple equivalent simultaneous offers tend to lead to better joint outcomes and better negotiator satisfaction. Furthermore, the strategy, when implemented in a software negotiation agent, can result in better counterpart agreement ratios, better individual and joint utilities, and better outcome satisfaction in an agent–human negotiation setting.[6]

It seems likely that both of these tactics—the voluntary sharing of interest information and the use of multiple equivalent simultaneous offers—can reduce mistrust between negotiating parties. Here, however, we emphasize the first of the two, and only in the context of an agent–human negotiation, where mistrust can be particularly problematic. Specifically, we're concerned with determining whether having a negotiating agent proactively communicate its issue priority to its human counterpart improves the negotiation outcome and improves the human's attitude toward the agent. Furthermore, if this tactic succeeds on both counts, the question arises as to whether it would still be effective for individuals who are particularly predisposed to mistrust others.

Given our belief that a good, trust-inducing tactic in a multi-issue negotiation is for one of the negotiators to volunteer its issue priority, and given the possibility that individuals who are especially predisposed to mistrust others might differ in the way in which they respond to the proactive disclosure of issue priority, we assessed participants on a variable that could capture such a predisposition. We did this using MACH-IV,[12] an established instrument that (in part) assesses cynicism and reluctance to trust others—that is, it measures *Machiavellianism*.

People with high scores on this test—*hi-Machs*—tend to endorse statements such as "Anyone who completely trusts anyone else is asking for trouble," but not those such as "Most people are basically good and kind." And, of course, if someone has this kind of view of people, he/she is unlikely to trust them. People with low scores—*lo-Machs*—tend to believe that humans are essentially good natured; lo-Machs thus typically take a more personal, empathetic approach in their interactions with others. They tend to be more trusting and more honest. Not surprisingly, research shows that lo-Machs and hi-Machs differ in both how they behave in and feel about negotiations in which they're involved.[12] See pp. 395–396 of Roy Lewicki and his colleagues' book, Negotiation, for more discussion of Machiavellianism and negotiation.[9]

Assessing the Effects of Proactive Communication

Two main groups of quantitative measures must be considered in evaluating a negotiation's effectiveness. The first group concerns the *economic outcomes*, which include the utilities derived from the agreement at both the individual and dyadic/joint level. Individual utility is assessed in terms of the extent to which the individual's outcome approaches an optimal value, whereas for dyads, outcomes are often measured by joint utility (such as the simple sum of individual utilities, or the agreement's distance from the set of Pareto-efficient solutions) attained by the negotiating parties. The second group of measures pertains to the *social-psychological outcomes* of a negotiation and includes measures of the negotiators' perceptions of the negotiation, the other party, and the self. Of these measures, outcome satisfaction and perception of the counterpart are among the best predictors of

a negotiator's willingness to negotiate with the same counterpart in the future.

As already discussed, hi-Mach individuals tend to be distrustful of other people, which means that they're less likely to see an agent's simultaneous-equivalent offers as an opportunity for a win–win outcome. However, if, as indicated above, proactively volunteering issue priorities provides an explicit signal that the agent is making an attempt to achieve a win–win outcome[10]—and especially if subsequent behavior confirms this interpretation—hi-Mach individuals might become less skeptical about the sincerity of their counterparts' offers. Thus, while still motivated by self-interest, they might be more likely to reach an agreement. In other words—and this is our main hypothesis—the use of such proactive communication should increase a hi-Mach's perception of the persuasiveness of the agent's offers, thus leading to better negotiation outcomes, whereas the same will be less true, if at all, for lo-Mach individuals.

This kind of proactive communication might also to lead to social-psychological or affective advantages in terms of personal perception—in this case, in terms of how people judge their counterparts. Because an agent's proffering of issue priorities can be seen as an attempt to share useful information aimed at obtaining mutual benefits, it can be interpreted as a sign of cooperativeness. Therefore, we expected that when an agent took the initiative to communicate its preferences to the human, that agent would be more likely to be perceived positively by the human—in particular, the human would more likely perceive the agent as considerate and flexible.

The Experiment

To investigate these questions, we compared the outcomes when humans with different levels of Machiavellianism negotiated with a "non-proactive" agent that provided only simultaneous-equivalent offers and a "proactive" agent that also volunteered its issue priority and invited the human to reciprocate.

Participants

We recruited MBA students enrolled in a negotiation dynamics class taught by one of the authors as participants. They were told that they would be playing the buyer role in an online negotiation, but they weren't told that the seller was a software agent. Participants were randomly assigned to negotiate with either a proactive or non-proactive agent.

A total of 54 participants (35 males and 19 females) completed the experiment. Their average age was 29. All participants had at least two years of work experience, with most (63 percent) being young professionals, and the others being more senior managers or directors. All reported (on five-point scales) high computer-usage experience (a mean of 4.31), moderate experience with online purchasing (3.76) and business decisions (2.93), and some experience with negotiation activities (2.04).

As an incentive to participate seriously, participants were told that they would receive tiered bonus course credit based on their negotiation performance, so that in addition to receiving five percent credit for their participation, the top 20 percent, the second 20 percent, and the third 20 percent of participants would receive an additional five percent, three percent, and one percent credit, respectively. The top negotiator would also receive a bottle of champagne worth approximately US$40.

Method

For this study, we adapted a multi-issue negotiation task originally developed from real-world supplier–buyer contract negotiations.[6] The task comprised four negotiation issues relating to the purchase of laptop computers: unit price, quantity, service level, and delivery terms. There were seven values for unit price, seven for quantity, four for service level, and four for delivery time (that is, 784 possible final agreements). Participants (buyers) and the agent (seller) had different weights over the issues and values, and were given a table that enabled them to determine the utility points any particular offer was worth. In game-theoretic terms, a non-zero-sum game was created that allowed for the possibility of an integrative, win–win outcome. To equate their bargaining power, buyers and seller were given the same 44 utility-point Best Alternative to a Negotiated Agreement (BATNA)—that is, the best bottom-line deal negotiators could obtain alone or with a third party if they were to fail to reach an agreement with the counterpart.

The experiment comprised three stages. In the first, *pre-negotiation* stage, participants read general instructions and were then briefed on the procedure. We then gave them a task sheet that described their role as buyers and indicated that their goal was to maximize their utility scores over the four issues for purchasing a new laptop computer model using the seller's website. Importantly, we instructed them to treat unit price as their top priority. Participants then took a quiz to make sure that they understood the task, after which, knowing their BATNA, they indicated the number of utility points (their target) with which they thought they would be satisfied. We then introduced them to the seller's website and the negotiation rules, which specified an alternating offer–counteroffer protocol with "final offer" termination rules.

Finally, participants completed a pre-negotiation questionnaire that assessed their business experience, computer usage experience, and so on.

In the second, *negotiation* stage, participants negotiated with the agent until they reached an agreement, or until negotiation ended with no agreement because one party rejected the other party's final offer. This phase had no time limit. In the third, *post-negotiation* stage, participants completed a questionnaire designed to assess their perceptions of the negotiation and the selling agent. Demographic information was also collected and—in an ostensibly separate and voluntary exercise—participants completed a "Negotiation-Self-Awareness" personality assessment questionnaire, the key part of which was Richard Christie and Florence Geis's 20-item Mach-IV test for assessing Machiavellian personality.[12] Before leaving, participants were asked to keep their experiences in the study confidential. The final debriefing, together with an announcement identifying the top performing participants, was conducted about two weeks after we collected all data from all participants.

Each participant negotiated with one of the two types of agent. For a participant negotiating with the proactive agent, the agent revealed its issue priority in the second round of the negotiation and invited the human to reciprocate. For a participant negotiating with the non-proactive agent—that is, the control condition—the agent simply presented its offers and reacted to the subject's counteroffers without ever communicating information regarding its issue priority.

In all other respects, the behavior of the two types of agents was the same. In both cases, the agent employed the same negotiation strategy to make and accept offers. Figure 1 shows the

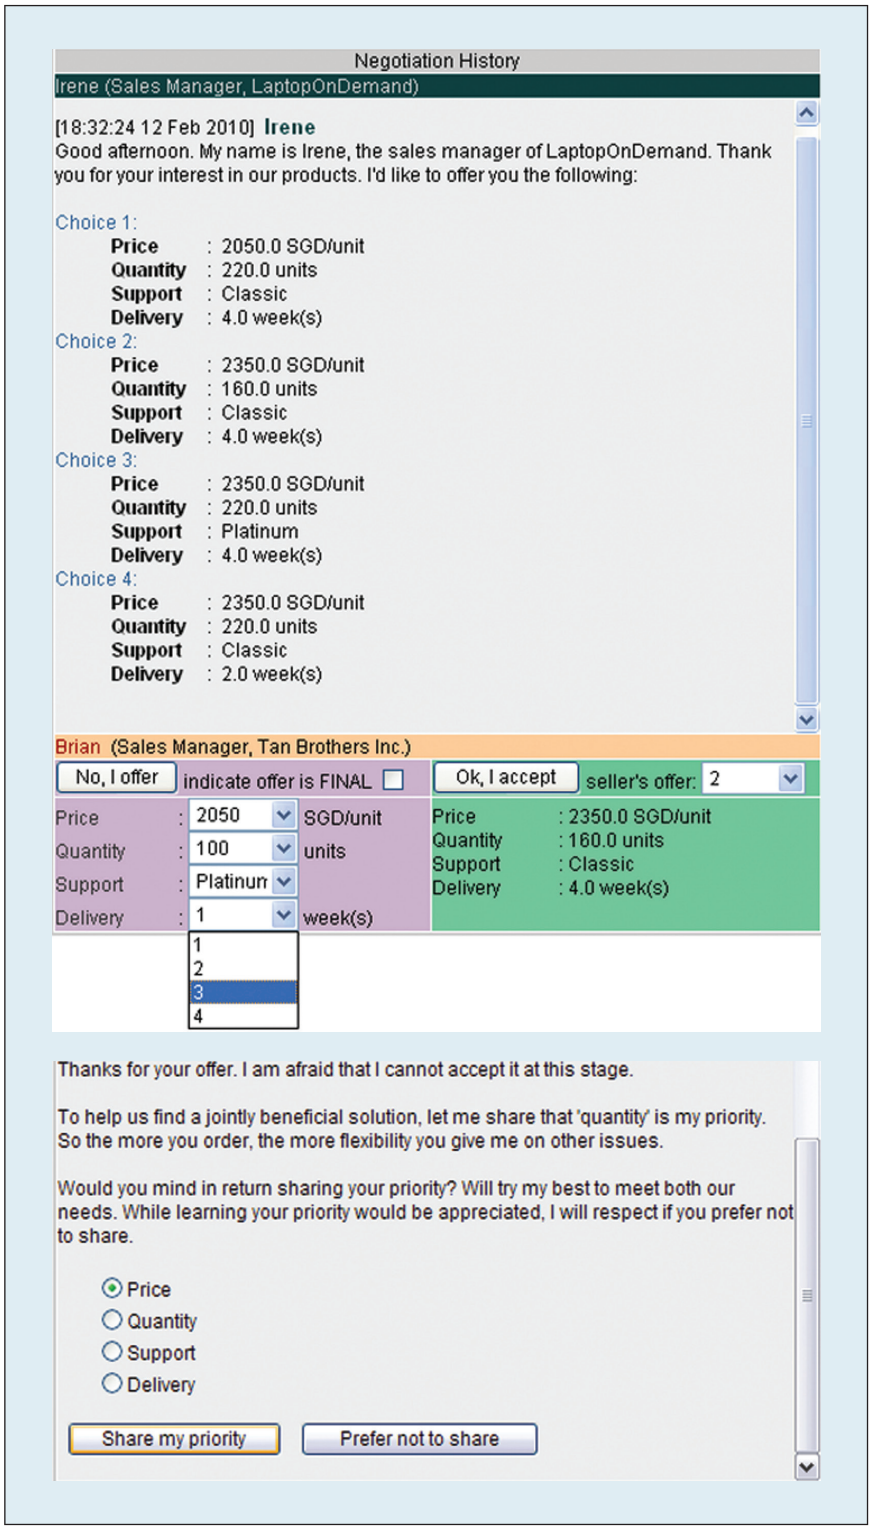

Figure 1. The agent–human negotiation interface, with the additional message presented by the proactive agent.

Table 1. Results of regression analyses on dependent measures.

Dependent measures	Predictor variables	β	SE	Wald/t	p-value
Agreements	Agent type	1.41	0.63	5.03	.03*
	MACH	–0.47	0.19	6.26	.01*
	Agent type x MACH	0.38	0.22	3.10	.08
Buyer utility	Agent type	7.32	3.05	2.40	.02*
	MACH	–2.11	.55	–3.83	.00**
	Agent type x MACH	2.06	.71	2.92	.01**
Seller utility	Agent type	3.15	1.88	1.67	.10
	MACH	–.70	.34	–2.06	.04*
	Agent type x MACH	.48	.44	1.09	.28
Joint utility	Agent type	10.47	4.22	2.48	.02*
	MACH	–2.80	.76	–3.70	.00**
	Agent type x MACH	2.54	.97	2.60	.01*
Outcome satisfaction	Agent type	.87	.47	.23	.07
	MACH	–.30	.08	–.70	.00**
	Agent type x MACH	.28	.11	.51	.01*
Perceived cooperativeness	Agent type	.44	.21	2.14	.04*
	MACH	.00	.04	.09	.92
	Agent type x MACH	–.01	.05	–.24	.81

*$p < 0.05$; **$p < 0.01$.

website interface, in which an example of the agent's four simultaneous equivalent offers is presented. Figure 1 also shows the message in which the proactive agent's issue priority—in this case, *quantity*—is presented (in the second round) and in which the human negotiator is invited (but not required) to reveal his or her own preferences.

Measurement

The dependent measures included

- the total number of agreements obtained (and the agreement ratio),
- the individual and joint utilities, and
- measures relating to social-psychological perspectives on the outcome and agent.

We used two items in the post-negotiation questionnaire to assess outcome satisfaction (how satisfied participants were with their utility scores) and three items to assess perceptions of the agent (the extent to which the agent was judged to be considerate, friendly, and flexible). The mean of the sum of these ratings constituted the collector variable—perceived cooperativeness.

Results

Control checks confirmed that there were no systematic relationships between agent type and participants' gender, past experience, and pre-negotiation target utilities. Also, a regression analysis revealed that the Machiavellian personality variable was unrelated to agent type, thus confirming that it can be assessed as a stable personality variable. After reverse coding items where appropriate, each subject's Mach-IV score was computed by summing his or her response values for all items.

Participants' mean score on the Machiavellian personality measure was 104.83 (with a standard deviation of 4.48, a minimum of 96, a maximum of 116, and a scale midpoint of 100), and the Shapiro-Wilk test indicated a normal distribution of scores. The construct reliability of the social-psychological dependent measures was adequate, with the items having acceptable reliability ($\alpha = .96$ for outcome satisfaction and $\alpha = .68$ for perceived cooperativeness).

Of the 54 agent–human dyads, 36 dyads obtained an agreement and 18 dyads did not. Of the 27 participants in the treatment (proactive agent) condition, 21 responded to the agent's invitation in the second round (see Figure 1), indicating, as they had been instructed, that unit price was their issue priority. Of the remaining six, three chose not to reciprocate, and three selected purchase quantity rather than unit price as their top priority.

We ran a logistic regression on agreements (because agreements was a dichotomous variable) and hierarchical regression analyses separately on buyer utility, seller utility, joint utility, outcome satisfaction, and perceived cooperativeness as the criteria, with agent type (a dummy coded binary variable of 0 = non-proactive vs. 1 = proactive) and

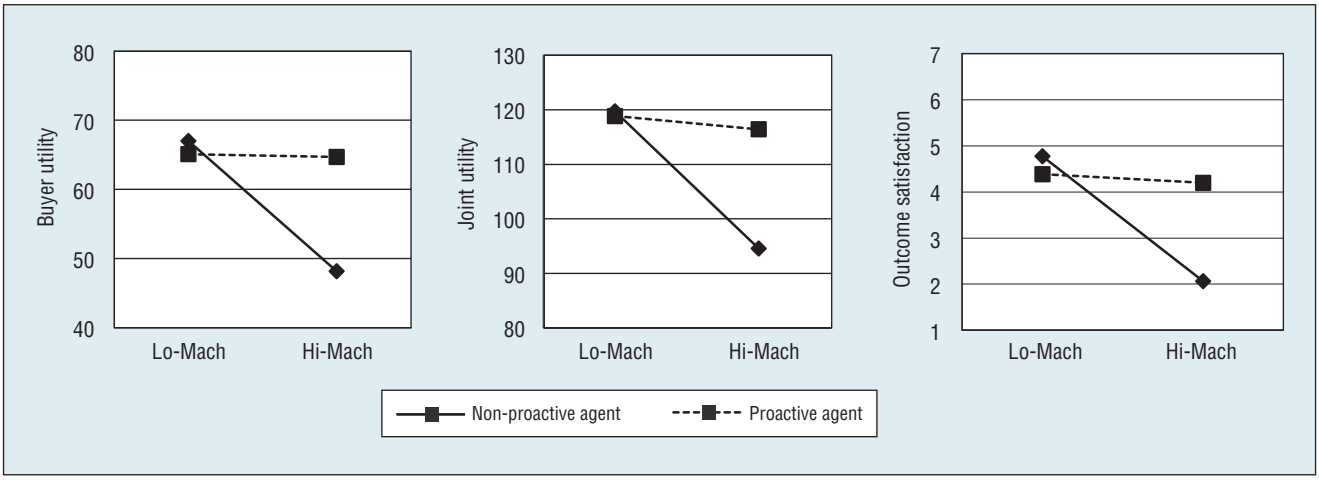

Figure 2. Buyer utility, joint utility, and outcome satisfaction as functions of agent type and Machiavellianism.

the centered personality variable, MACH, as predictors in each case. Table 1 shows the results of all of these analyses.

Agreements

Overall, out of 27 cases, negotiating with the proactive agent led to 22 agreements (an 81.5 percent agreement ratio), whereas negotiating with the non-proactive agent led to 14 agreements (a 51.9 percent agreement ratio). This main effect was significant. With regard to the influence of the Machiavellian personality variable, results also showed that MACH had a significant negative relationship with the agreements. Agent type and MACH were found to have a marginally significant interaction.

Other Dependent Measures

We found similar patterns for the other dependent measures of buyer utility, joint utility, and outcome satisfaction, with significant main effects of agent type and MACH, as well as an agent type by MACH interaction. For perceived cooperativeness, results indicated a main effect of agent type only, with the proactive agent being perceived as more cooperative than the non-proactive agent. Neither the main effect of MACH nor its interaction with agent type was significant. We also analyzed two other dyadic outcome measures—the distance to Pareto efficient frontier and the distance to Nash solution—and the results showed the same pattern as for joint utility.

We conducted further analyses to examine the effects of agent type on hi- versus lo-Mach individuals using simple slope analysis. The scores of both hi- and lo-Machs were centered half a standard deviation above and below their respective means. The equivalent ranges of Mach-IV scores were 96–102 for lo-Machs ($n = 17$ and mean = 99.7), and 107–116 for hi-Machs ($n = 19$ and mean = 109.8). Results indicated that, for hi-Machs, the proactive agent led to significantly better negotiation outcomes than did the non-proactive agent in terms of agreements (7 out of 10 and 2 out of 9, respectively); buyer utility ($\beta = 15.52$, $SE = 5.38$, $t = 2.83$, and $p = 0.01$); joint utility ($\beta = 18.73$, $SE = 7.65$, $t = 2.45$, and $p = 0.03$); and outcome satisfaction ($\beta = 2.26$, $SE = 0.86$, $t = 2.64$, and $p = 0.02$). For lo-Machs, there were no such differences. Figure 2 shows the interaction effects.

Discussion

The data clearly bear out our predictions concerning the effects of a negotiation agent proactively volunteering information about its issue priority (and gently inviting the counterpart to do the same). Specifically, relative to trusting individuals, less trusting, more skeptical (hi-Mach) individuals achieved better negotiation outcomes when the agent volunteered information about its issue priority than when it didn't. This was evidenced by the significant interaction between agent type and Machiavellianism wherein more agreements, better buyer and joint utilities, and greater buyer satisfaction were achieved for hi-Machs but not for lo-Machs when negotiating with the proactive agent. From the social-perception perspective, the results also indicated that participants viewed the proactive agent more positively than the non-proactive agent.

By showing that the simple gesture of offering cooperative information exchange can have beneficial economic and affective consequences, these findings demonstrate the value of taking the dilemmas of honesty and trust[9] into account when seeking an integrative win–win outcome. Clearly, adopting a proactive stance and volunteering issue priority is a good move. However, our results indicate that, from an economic perspective, the benefits are sensitive to personality factors in a way that appears not to be true from the social-psychological (affective) perspective. In our admittedly limited

sample, an economic advantage was observed only for hi-Mach individuals—that is, for individuals who tended not to trust others. In contrast, the affective advantage was independent of both personality differences and economic outcomes: participants simply felt better about the proactive agent than about the non-proactive agent, even though information gathered during the debriefing phase indicated that most participants realized that they had been negotiating with a software agent.

The fact most of our participants (21 of 27) responded to the agent's invitation to reciprocate by revealing their own issue priority has an interesting implication for designing more cognitively and socially intelligent negotiating agents, particularly when such agents have to negotiate with skeptical, untrusting counterparts. This evidence of cooperativeness by the human negotiator suggests that the simple mechanism of volunteering issue priorities has the capacity to unlock a negotiation's win–win potential.

Whereas earlier negotiation agent research has focused on techniques for inferring preferences,[2-4] we've explored a direct rather than an inferential way of foregrounding information about preference tradeoffs that nevertheless increases the agreement ratio and joint utility of agreements. Our results show that a software negotiation agent volunteering its issue priorities and inviting the counterpart to reciprocate has positive economic consequences, at least for untrusting cynics. At the same time, it has positive affective consequences vis à vis attitudes towards the agent, even when there are no immediate economic gains (as was the case for lo-Machs). It's also possible that developing a positive affect toward the counterpart would have economic benefits in negotiations with the same agent in the future. It's thus apparent, especially given the low potential risks and the high potential rewards, that it makes sense to design negotiation agents to initiate cooperative win–win moves by volunteering issue priorities and inviting reciprocation.

Our research also provides new insights into the relation between Machiavellianism and negotiation behaviors. Past research indicates that hi-Machs get better utilities than lo-Machs, especially when they negotiate with lo-Machs.[13,14] However, we found that hi-Machs performed *worse* than lo-Machs when they interacted with a non-proactive agent, a finding that's similar to that of William Fry,[13] who found that in a condition without visual access, hi-Mach/hi-Mach dyads performed worse than hi-Mach/lo-Mach and lo-Mach/lo-Mach dyads. To the best of our knowledge, ours represents the first time an attempt has been made to influence hi-Machs to cooperate by encouraging them to see that it's to their advantage to do so. This suggests that even though hi-Machs are generally more cynical and self-interested than others, it's possible to engineer situations that will lead them to a cooperative approach to negotiation.

As the above discussion implies, our conclusions are limited to specific negotiation contexts—that is, those involving multiple issues with preference tradeoffs in which integrative, win–win outcomes are at least a possibility and the negotiating agent actually intends to encourage win–win seeking behaviors (that is, it volunteers its real issue priority early on and behaves consistently in subsequent offers). It's unrealistic to expect a negotiator to trust an agent—human or otherwise—whose behavior is inconsistent, as such inconsistency would generate low expectations about a negotiation's economic or affective consequences.

Also, the reduction of mistrust in hi-Machs is only an inference from our data; we didn't directly measure the extent to which participants trusted (or mistrusted) their negotiation counterparts. However, we believe that it's the most reasonable and parsimonious explanation of the interaction that we observed. As Figure 2 clearly shows, the main thing to be explained is why the outcomes for hi-Machs was so much better when the agent made its cooperative gesture (and essentially the same as for lo-Machs in both conditions). The best explanation is surely that there was something about the gesture that removed the disadvantage that hi-Machs otherwise had—specifically the agent's signal of cooperativeness reinforced by its subsequent behavior. In other words, their trust in it increased.

Finally, there's an important issue pertaining to the agent's nature and the fact that, upon debriefing, most participants in our study said they believed they had negotiated with a computational agent. Of course, we have no idea at what point participants came to this realization. Was it only after the negotiation, upon reflection of what had happened? Or was it during the negotiation—and if so, was it early or late in the process? Was there anything in particular that tipped them off? It might be fruitful to explore such questions, as well as to examine the consequences of trying to manipulate participants' beliefs about the agent's nature—human or computational. However, one problem with trying to do this is that even if we tell participants that they're negotiating with a human, they might not believe it. Indeed, informal reports from our participants suggest that even lo-Machs might well be suspicious. On the other hand, we could argue that

because people tend to treat machines as if they were humans,[15] such a manipulation would be unlikely to affect the outcome. These and similar questions will have to await future research. ■

References

1. P.B. Linhart, R. Radner, and M.A. Satterthwaite, *Bargaining with Incomplete Information*, Academic Press, 1992.
2. D. Zeng and K. Sycara, "Bayesian Learning in Negotiation," *Int'l J. Human–Computer Studies*, vol. 48, no. 1, 1998, pp. 125–141.
3. P. Faratin, C. Sierra, and N.R. Jennings, "Using Similarity Criteria to Make Negotiation Tradeoffs in Automated Negotiations," *Artificial Intelligence*, vol. 142, no. 2, 2002, pp. 205–237.
4. R. Ros and C. Sierra, "A Negotiation Meta Strategy Combining Trade-Off and Concession Moves," *Autonomous Agents and Multi-Agent Systems*, vol. 12, no. 2, 2006, pp. 163–181.
5. R. Lin and S. Kraus, "Can Automated Agents Proficiently Negotiate with Humans?" *Comm. ACM*, vol. 53, no. 1, 2010, pp. 78–88.
6. Y. Yang, S. Singhal, and Y. Xu, "Alternate Strategies for a Win–Win Seeking Agent in Agent–Human Negotiations," *J. Management Information Systems*, vol. 29, no. 3, Winter 2012–13, pp. 223–255.
7. C.E. Naquin and G.D. Paulson, "Online Bargaining and Interpersonal Trust," *J. Applied Psychology*, vol. 88, no. 1, 2003, pp. 113–120.
8. H.H. Kelley, "A Classroom Study of the Dilemmas in Interpersonal Negotiation," *Strategic Interaction and Conflict: Original Papers and Discussion*, K. Archibald ed., Inst. of Int'l Studies, Berkeley, 1966, pp. 49–73.
9. R.J. Lewicki, D.M. Saunders, and B. Barry, *Negotiation*, 5th ed., McGraw-Hill, 2006.
10. H. Falcão, *Value Negotiation: How to Finally Get the Win–win Right*, Prentice Hall, 2010.
11. V. Medvec et al., "Choice and Achievement at the Bargaining Table: The Distributive, Integrative, and Interpersonal Advantages of Making Multiple Equivalent Simultaneous Offers," *Proc. 18th Ann. Conf. Int'l Assoc. Conflict Management* (IACM), 2005; http://dx.doi.org/10.2139/ssrn.732665.
12. R. Christie and F.L. Geis, eds., *Studies in Machiavellianism*, Academic Press, 1970.
13. W.R. Fry, "The Effect of Dyad Machiavellianism and Visual Access on Integrative Bargaining Outcomes," *Personality and Social Psychology Bull.*, vol. 11, no. 1, 1985, pp. 51–62.
14. A. Gunnthorsdottir, K. McCabe, and V. Smith, "Using the Machiavellianism Instrument to Predict Trustworthiness in a Bargaining Game," *J. Economic Psychology*, vol. 23, no. 1, 2002, pp. 49–66.
15. B. Reeves and C. Nass, *The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*, Univ. Chicago Press, 1996.

THE AUTHORS

Yinping Yang is an independent investigator, scientist, and program manager at the Institute of High Performance Computing, A*STAR, Singapore, and an adjunct faculty of information systems at the Singapore Management University. Her research interests include intelligent negotiation agents, e-commerce, social media, and cloud computing. Yang has a PhD in information systems from the National University of Singapore. Contact her at yangyp@ihpc.a-star.edu.sg.

Horacio Falcão is a senior affiliate professor of decision sciences at INSEAD and a cofounder of Pluris Pte. Ltd., a negotiation consulting company. His research interests include strategies in win-win negotiation, cross-cultural negotiation dynamics, and organizational change. Falcão has a Master in Business Administration from INSEAD. Contact him at horacio.falcao@insead.edu.

Nuno Delicado is a founding partner of Pluris Pte. Ltd., a negotiation consulting company. He conducts negotiation and mediation training, coaching, facilitation, and consulting to both the private and public sectors. Delicado has an MBA with distinction from INSEAD. Contact him at nuno.delicado@plurisvalue.com.

Andrew Ortony is a professor emeritus at Northwestern University and past director of the Programme in Computational Social Cognition at the Institute of High Performance Computing A*STAR, Singapore. His recent research interests include computational models of affect and emotion and their relationship to personality. Ortony has a PhD in computing science from the University of London's Imperial College of Science and Technology. Contact him at ortony@northwestern.edu.