

Caretaker: A Social Game for Studying Trust Dynamics

Nicholas Violi*, Jennifer Golbeck*, Kan-Leung Cheng*, and Ugur Kuter†

*University of Maryland; {nvioli, golbeck, klcheng}@cs.umd.edu

†Smart Information Flow Technologies; ukuter@sift.info

Work done while at the University of Maryland

Abstract—We present a new strategy for studying trust determination, especially in the context of dynamic trust. We created a game in which the players start with neutral trust for each other, have assigned roles (“good” or “bad”) and an incentive to deduce the roles of the other players, and are encouraged to take part in social and economic interactions with each other, thereby gathering data with which to make trust determinations. By running user tests, we show that a game with these components generates useful data with which the players’ decisions regarding trust for each other can be observed and better understood. We show that players are more likely to increase their self-reported trust for each other after engaging in an economic transaction, and, surprisingly, their trust for each other increases by a larger margin when the transaction was asymmetric. We also present an analysis of the chat that the players participated in, via an in-game chat system. We show that the two most popular topics of discussion are trade requests and game administration, and suggest possible future work in comparing the players’ chat actions to their trust for each other.

Index Terms—games, trust, dynamic trust, social games, social strategy

I. INTRODUCTION

Trust is one of the most important contributors to successful social interactions. It supports our decisions to take risks based on the actions or advice of others. Understanding how trust is established between people and how it evolves based on their respective actions is a challenging task, in part because so many complex factors contribute to the development of trust: our history of interaction with that person, the success or failure of those interactions, our trust in the friends and associates of the person, the similarity between ourselves and that person (on any number of factors), and so on. Even trust models using artificial agents, which have much simpler social lives than people, have included many of these complex factors [1].

Furthermore, trust between two (or more) parties is not a fixed quantity; it is constantly changing in response to the actions of and interactions between them. In turn, the parties’ actions are affected by their evolving trust for their counterparts, so one party’s trust for other parties, and his or her actions towards them are in a constant state of flux, and the two are continually affecting each other in complicated ways. The mechanisms and triggers for the dynamics of changing trust are important to understanding behavior in social situations, and yet to date they are not well understood.

There are many problems where an understanding of social strategies and dynamic trust decisions is important, such as economic negotiations and markets, professional and social communities, political diplomacy, as well as technology-mediated situations, both cooperative and competitive. Mathematical models for dynamic trust have been proposed by [2] and [3], but there has been relatively little research into strategies for empirical testing of such theories.

To study social strategies, trust, and the dynamics of trust, we have developed a game with features that allow us to measure and quantify behavior, its impact on trust, and vice versa. We required the game have the following features:

- Trust is initially unknown.
- There must be economic and social interactions to serve as bases for establishing a history that will influence trust.
- Players must have an incentive to investigate and establish trust in others.

By building a game that satisfies these three requirements, we were able to observe how players formulate decisions on whom they trust based on their actions and those of their opponents, how their actions change according to these trust determinations, how their trust evolves based on their and their opponents’ evolving and changing behavior, and so on.

The game provides a framework for studying trust and understanding how trust affects social strategies. In our first round of human testing, we were able to observe some elementary factors that contribute to or are correlated with changes in trust. We will show that it is possible to observe changes in players’ trust for each other in response to the following events:

- Any interaction
- Economic interactions
- Social interactions
- Strategic actions

In this paper, we begin by presenting related work on using games for studying social interaction. This is followed by a description of the game framework we developed, a description of our current implementation, and results of user tests that illustrate the type of data we can obtain through our studies.

II. RELATED WORK

Before a decade ago, research interest in games focused primarily on studying them for their own intrinsic qualities,

in the field known as game theory, or as models for other phenomena, as in economics. Since then, partially spurred by a 1999 paper by Fehr, et al.[4] which sought to explain certain observations regarding social interactions between players of games, they have increasingly been used to study many facets of the social interaction that arises between agents playing games.

Perhaps the canonical example of such a study concerns a game called “Colored Trails”[5][6]. This game has been used to study cooperation and socialization between players (in either cooperative or antagonistic positions); among the major items researched by the Harvard group is the idea that by cooperating, players can perform better than when using selfish methods suggested by traditional game theory. Many other games have been used to study the phenomenon of cooperative advantage[7], and games have been used to study many other facets of socialization and trust between players of games[8], such as considering social preferences of opponents[9][10], and optimal strategies for forming and breaking alliances and subgroups[11].

Games have also been created to study certain facets of interaction in a business setting, including supply chain management[12], implications of distributed or colocated teams[13], and optimal responses to requests for proposal[14].

Finally, there has been much interest in “games with a purpose”[15]. Here, the goal is not to study the social interactions of the players, but to create an enjoyable game which has the side-effect of producing data that is both useful and computationally difficult to create without human intervention. Such games have been used to gather data on image-recognition tasks[16], to create metadata for photos for accessible internet-browsing[17], to gather common-sense facts[18], to label images[19], to create metadata and tags for music[20], and to group objects into collections[21].

III. GAME BASICS

Our game, Caretaker, is designed to model a transportation network in a setting where the state of the terrain is unknown (e.g. perhaps there are terrorists controlling access to some places, or there has been a natural disaster rendering some of the roads impassable). The players in the game represent different factions (e.g. government forces, non-governmental organizations, insurgents) who all have specific goals in the region, and whose motives and loyalties are initially unknown to the other factions. Each player must rely on the other players for help in achieving his or her own goal, and yet he or she must attempt to determine which players share his or her loyalty and which do not, to aid the former in fulfilling their goals, and to prevent the latter from fulfilling theirs.

Within this framework, the game is designed to be flexible in its design to support controlled testing of different principles. At its core, the game has a basic set of rules and principles that supports the analysis of actions and dynamic trust determination, as described in the previous section. The game’s basic requirements are as follows:

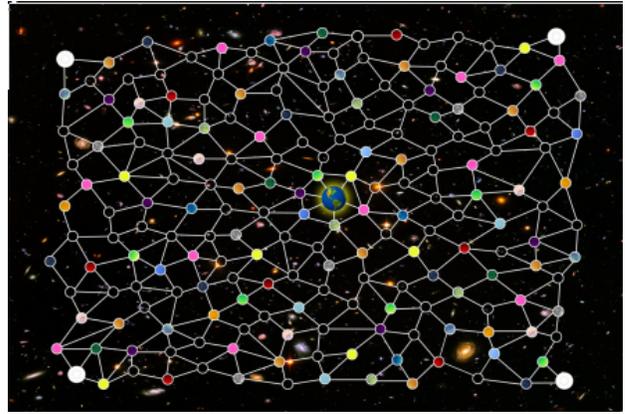


Fig. 1: The currently used Caretaker gameboard

- 1) The game board has a network structure. Players move from node to node from the outside, toward a goal at the center, crossing one edge each round. The board we’re currently using is shown in Fig. 1.
- 2) Resources (e.g. money, fuel, etc.) are required to cross each edge. These are distributed each round, but the distribution of resources is sparse enough that players will often need to negotiate, trade, and exchange resources in order to make a move. Without the correct combination of resources, the player must stay in the same place until the correct resources are obtained.
- 3) At the beginning of the game, players learn if they are a “good” or “bad” player. This information is not shared with the other players.
- 4) In order to win, the good players need to reach the center goal before the bad player(s).
- 5) Through negotiation and exchanges, players try to identify which of their opponents are good and which are bad.

As stated above, it was important to us that players begin the game with neutral trust toward all the other players, so that we could be sure that any trust or distrust that was developed was solely related to the players’ interactions within the game (which could be recorded and later played back or analyzed). This requirement is met by divulging the players’ own roles (good or bad), but leaving their opponents’ roles a mystery. By making the players’ success or failure in the game contingent not only on the speed at which they reach their goal, but on the speed of their entire team, we incentivize the players to try to discern or deduce the alignments of the other three players (i.e. to determine their trust for them, and to continually refine their determination) over the course of the game. Finally, by requiring that the resource distribution be sparse, and by providing spaces to negotiate and trade for resources (the trade system), as well as a freeform social area in which the players can discuss any relevant topics (the chat system), we provide both the facility and the incentive to participate in social and economic interactions, thus providing the players with social data on which to base their trust determinations. Thus, this setup ensures that the three features described in the Introduction are met.

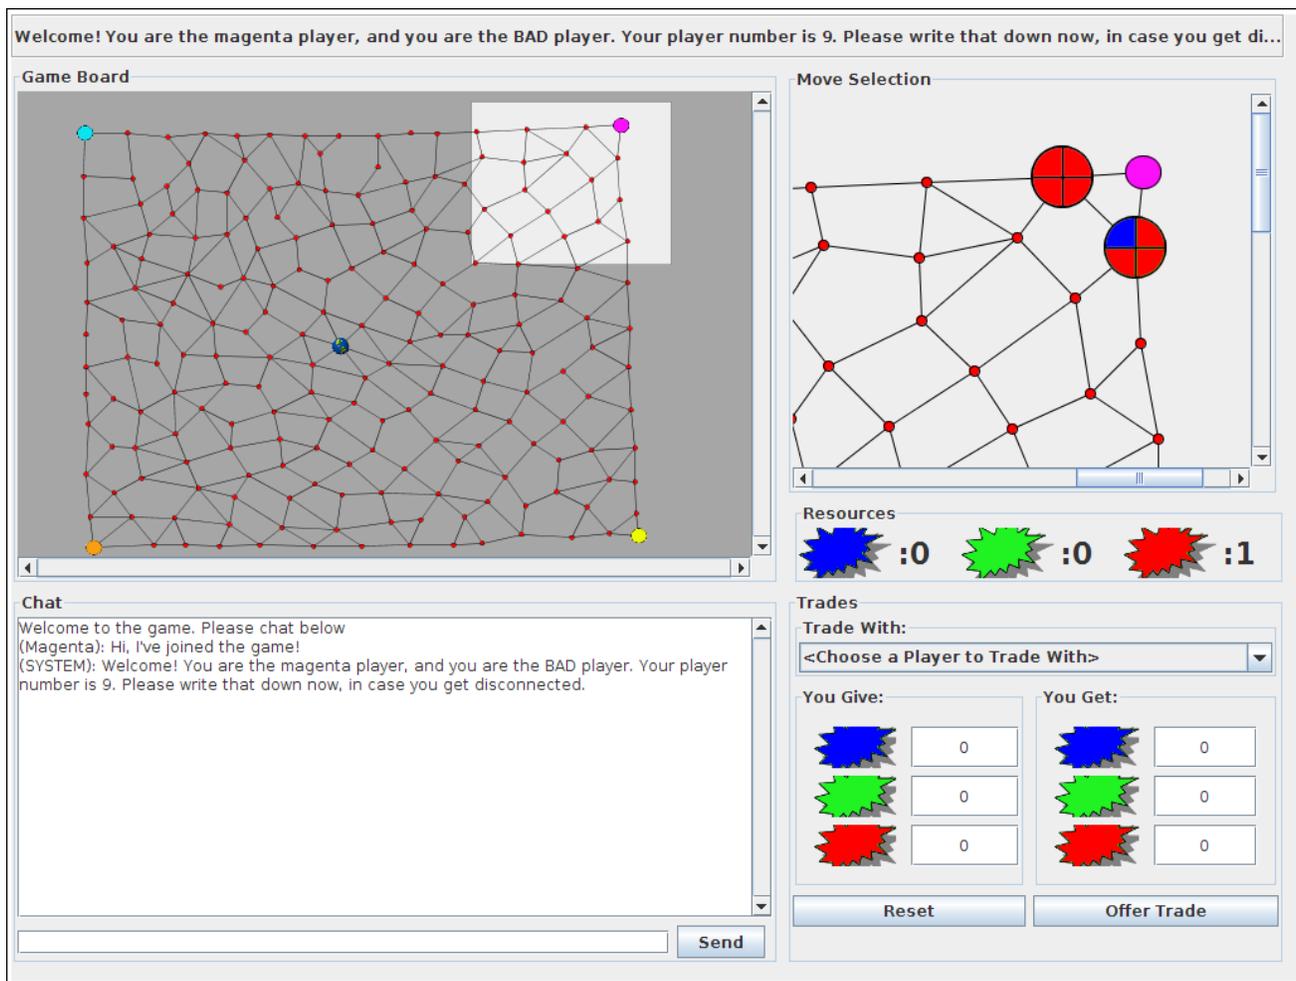


Fig. 2: The Caretaker applet window

The most basic data that is available to a player regarding his or her opponents is the strategic decisions they make. The game involves moving through a network toward a goal, and in the current implementation this goal is common for all players, though it need not be. The location of all players at any point during the game is public information, and (in the current implementation) their paths to the goal all have the same minimum length, and roughly the same branching factor. Since all players on the good side must reach the goal before the bad player does, a player who moves quickly ahead of the other players may be seen as unwilling to share resources, and thus may arouse suspicion that he or she might be on the bad side.

The chat system enables players to send arbitrary text to each other, such as proposing trades, or making accusations about the other players. This may be a valuable way for players to determine whom to trust: how talkative the other players are, how fast they want to move, how greedy or generous they are with resource requests or offers, and whom each player professes to suspect may be on the "bad" team can all provide the players with social and strategic information

about their opponents. These data can be used by the players to develop opinions on the alignment of their counterparts, or to strengthen, weaken, or otherwise change their opinions of whom they trust and distrust. The last factor, the accusations that players make toward each other, is quite similar to the social data which dominates the game "Mafia"[22], [23].

The trade system allows players to propose trades to their counterparts. Because of the resource sparsity, this is almost guaranteed to be a necessary part of the game for a player who wishes to be successful. Through this structured system of economic transactions, the players are given another venue for gathering social and strategic information about their counterparts. A player's willingness to trade, their generosity or selfishness in those trades, their flexibility, and the ratio or the number of trades they propose to the number of trades they accept all provide information which the other players can use to factor into their decision of whether or not to trust that player. Also, because winning or losing depends on supporting others on the player's team, while limiting the resources given to the opposing players, there is an incentive for increasing interaction to learn whom to trust and support.

Finally, the trade system serves as a key component in the functioning of the game, as once the players have made strong trust determinations, the trade system allows them to share scarce resources with those they trust, and to withhold them from those they distrust.

IV. IMPLEMENTATION

To illustrate these principles, we will describe our current instantiation of the game. Caretaker is currently set up as a 4-person game, and the entire user interface is shown in Fig. 2. Three players are good and one is bad. They begin on the outside corners of the rectangular board and move toward the center. For the good team to win, all three must arrive before (not at the same time as) the bad player. If the bad player reaches the center at the same time as or ahead of even one good player, he or she wins.

The chat system is implemented as a panel in the applet's lower-left corner that facilitates public chat between the players, as well as notifications from the game. Each time a player enters text into the chat box, it appears on all other players' screens, identified by their individual color. Anonymous statements and private chat are not implemented in this version of the game.

The trade system is implemented as follows: using the panel in the applet's lower-right corner, players may offer each other trades at any time before they have locked their move (see below). Upon receiving a trade proposal, a player can accept or reject it; if they reject it, they may offer a rebuttal. Players may offer any number of resources to any other player, provided that they possess enough to in fact complete the trade. They may ask for any number of resources in return. Note that the trade system supports trades of equal numbers of resources, trades of unequal numbers, and degenerate trades where the number of resources being offered or being requested is zero (enabling the players to "give" or "request" resources for nothing in return). Trades are strictly private in this version of the game: a trade must be proposed by one player to exactly one other player, and the trade is only visible to those two players. Finally, players are only allowed to have one trade in progress at a time.

There are three types of resources: "red", "green", and "blue". It costs four resources to cross an edge and make a move, except for the edges connected to the goal, which cost eight resources to cross. The specific resources required to make a move are generated randomly, and do not change between turns. The costs of all available moves are displayed in the "Move Selection" panel in the applet's upper-right corner. At the beginning of each turn, the players are given a random assortment of between zero and five resources, they then begin discussions using the chat and the trade panels. When each player is finished negotiating, he or she "locks" his or her move (which is allowed to be the stationary move) using the Move Selection panel. Players may not trade when they have locked their move, but they may unlock and lock again as many times as they wish. When all four players have locked

their move at the same time, all moves happen simultaneously and the next round begins.

The image shows a window titled "Survey" with a close button in the top right. The text inside reads: "Please take a moment to rate your trust for your opponents. 100% means you are sure this opponent is 'good'; and 0% means you are sure this opponent is 'bad'." Below this are three identical rating scales for "Cyan", "Magenta", and "Yellow". Each scale has a horizontal line with "Bad" on the left, "Good" on the right, and "Unsure" in the center. A small shield icon is positioned on the line, roughly in the middle. At the bottom of the window is a "Submit" button.

Fig. 3: Survey form presented to all players at the end of each round

Finally, at the conclusion of every round, we presented all players with a short form which required them to rate their three opponents based on whether they thought they were on the good or bad team, and how confident they were in that opinion. An example of this survey can be seen in Figure 3. This information was kept completely private with respect to the other players, and was crucial to our analysis of the game data. The players' ratings for each other, recorded after every game round, were used as signals for their trust for each other. Thus, by analyzing the players' ratings for each other, and especially by investigating the changes in those ratings, we were able to observe what actions and interactions within a round had the biggest effect on trust.

Our implementation of the game also contains a logging system which records the following information:

- When the players lock and unlock potential moves (along with what move is selected),
- When all players have locked their moves, and therefore the moves progress,
- How many resources the players receive at the start of each turn,
- The cost of all moves available to each player,
- Information about when a trade is offered, accepted, rejected, or rescinded,
- All chat actions by all players, and
- The results of the surveys presented at the end of each round.

See the next section for more details.

V. DATA GATHERED

We held open sessions of gameplay to gather data. Over the course of two months, we gathered data on 19 games from friends, graduate students, and undergraduates. This represents data on 76 players; 55 of them good and 19 of them bad. The shortest path to the goal from any of the four starting positions is ten moves, but because of the purposeful scarcity of resources, it would take an extraordinarily lucky player (or perhaps very naïve opponents) to finish in only ten turns. The total number of rounds recorded was 381, for an average of 20 rounds per game.

In each round of gameplay, when a trade was offered, we recorded the offering player, the player the trade was offered to, and the amounts and types of resources in the proposed trade. When a trade was responded to, either by acceptance or rejection, or if the offering player rescinded the offer, this was recorded. There were 380 trades proposed in the gathered data, or approximately 1 per round. Of these, 203 (53%) were accepted and 177 (47%) were rejected or rescinded.

We also recorded the responses to the survey questions, resulting in 1493 ratings (due to players' connectivity problems during some of the games, a small number of survey results were not recorded). Note that survey responses completed by the bad player are largely meaningless since they know with 100% certainty the roles of all players. Bad players responded to this fact in different ways: rating all opponents as 100% good, leaving the sliders at their default value of 50%, choosing their answers randomly, or using some other scheme. Therefore bad players' survey data are not reported for the remainder of the paper.

Finally, we recorded all chat actions. There were a total of 6131 chats sent, or 16 per round, or 4 per player per round.

VI. RESULTS

From the data described in the previous section, a number of interactions were found to have an effect on trust between players. The largest such effects concerned trades between two players, and are summarized in Table I. Each occurrence in this table represents an event involving a particular player who was on the good team. In the case of trades, each individual trade may be listed twice, since the two players participating in the trade each reported their own trust values at the end of the round. If a player participated in more than one trade with the same person in a single round, the survey data is attributed to all trades, since it is impossible to determine which (if any) trade affected the survey result.

A. Chat Log Analysis

We logged chat conversations throughout the games because this provides insights into how players are communicating and negotiating. For this analysis, we studied

the chat logs independent of the quantitative data, with a goal of understanding what players were discussing and with what frequency.

Among all the games, there were 5,933 chat messages sent after log cleanup to eliminate blank messages, system codes, and the like. To understand the content of the messages, we developed a code book using an emergent coding technique [24] with two trained expert coders. The codes developed were as follows:

- Game Administration - These are discussions about the game function including locking moves, asking about locks, questions about game processes and procedures, identifying who is who in the game, etc.
- Greetings - Chat is social, so the logs included many greetings and friendly interactions like "Hi", "how are you?", etc. which were categorized together.
- Trash Talk - Different from greetings, trash talk was playful banter about winning and losing, with statements like "You're going down", etc.
- Requests / offers of trades and resources - Trades and their negotiations are one of the most critical parts of the game, and the motivation for having a chat service in the first place. Example statements included "I need 2 greens", "I have extra blue. Does anyone need them?", and "I'll give you 1 red for 2 blue."
- Theories or Accusations about who is good/bad, Statements about trust - As the game goes on, players begin to guess at who is good or bad and share that with the group. Toward the end of the game, this becomes a particularly important part of trading strategy. Examples include "Magenta is totally the bad guy" and "I'm good!".
- Resource complaints / statements / questions - Even when players are not discussing trades, they often have things to say about their resources. It may be a complaint about not having enough or gripe with what they were randomly awarded on the current move. Example messages are "I didn't get anything on that turn!" "The resource allocation is painful", "I have 12 reds and nothing else!", and "Did you get any resources that turn?"
- Non-Game talk - Some players have conversations in chat that are totally unrelated to the game. This may be general chit chat about unrelated topics, an exchange of contact information, and the like.
- Exclamations - Outside of any particular conversation thread, users would interject expressions like "lol", "crap!", and emoticons.
- Move talk - Understanding the state of other players is important to the game. Knowing who can move, who has been stuck, and why are all critical pieces of information. Examples of this type of message include "Who can move?", "Are you stuck?", and "How long since your last move?"

| Event | Occurrences | Trust increases | Amt of increase | Trust decreases | Amt of decrease | Trust unchanged |
|-------------------------------------|-------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| All rounds | 1386 | 239 | 4327 | 232 | 3812 | 915 |
| No trade completed | 962 | 158 | 2584 | 167 | 2751 | 637 |
| Favorable completed trades | 111 | 25 | 653 | 27 | 511 | 59 |
| Unfavorable completed trades | 111 | 31 | 560 | 18 | 266 | 62 |
| Mutually favorable completed trades | 202 | 25 | 530 | 20 | 284 | 157 |
| Rejected (or rescinded) trades | 177 | 26 | 588 | 33 | 471 | 112 |

TABLE I: Results

When coding messages, we kept conversation threads together where possible, so messages were coded in the context of the ongoing discussion.

Two coders classified each message and we achieved 96.7% inter-coder agreement. Only messages on which the coders agreed were used in this analysis, a total of 5,742.

The most common messages, not surprisingly, were trade requests and negotiations which were 39% of all messages sent. Since this type of interaction is critical to the game and requires a lot of communication, we expected this would be a popular category. This is a positive sign for the development of the game, since the trading system was created as a mechanism for players to develop trust and theories about who is on which side. The fact that players are actively communicating around this topic suggests the dynamics of the game are playing out as desired and helping users make decisions.

Game administration, at 32%, was the next most common type of message. Most of this was players announcing that they had locked in their moves, a necessary step for the game to progress. Many of these messages could be eliminated by changes in the user interface of the game that would expose some of the information people are sharing, but there is no reason this would impact trust and thus we do not see it as having bearing on these results.

The remaining types of messages all together were used less frequently than either of the two most popular categories. Expressing theories about who is good or bad and who is trusted was the most popular of the remaining categories. There were 387 of these messages, and as future work we would be interested to analyze how statements about these theories from players correlate with the actual scores they assign in the trust survey.

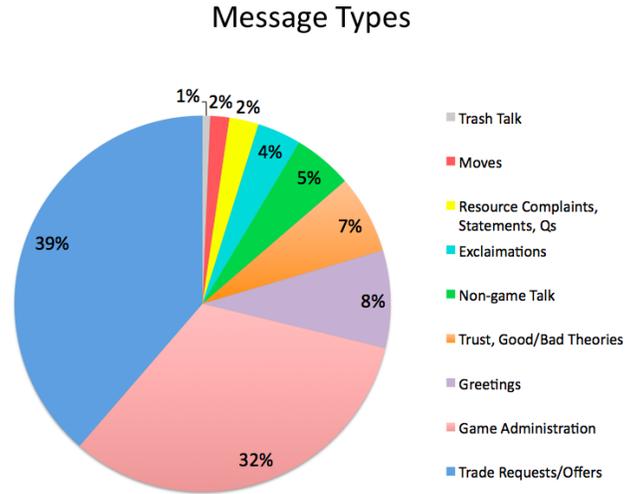


Fig. 4: Message types used in chat transcripts.

VII. DISCUSSION

We can use the data gathered in the gameplay to learn about what factors affect players' trust for each other, and, conversely, what actions are affected by their current trust level. For instance, investigating the rounds in which a trade occurred, we find the following result: in a round where no trade was completed, the trust of one player for another decreased by an average of 0.17 points $((2584 - 2751)/962)$, however, in the rounds corresponding to the 424 completed trades, trust for the other player involved in the trade increased by an average of 1.6 points $((653 + 560 + 530 - 511 - 266 - 284)/(111 + 111 + 202))$. Thus we can see that after a round in which two players participate in a mutually-agreed-upon trade, their trust for each other will increase substantially, and in a round in which they do not trade, it will decrease slightly.

Furthermore, if we look at only those trades which were imbalanced (a trade in which the resources being given and those being received are not equal), we find a stronger result: an average increase in trust of 2.0 points per trade $((653 + 560 - 511 - 266)/(111 + 111))$. Perhaps even more surprisingly, this effect is strongest when the trade is unfavorable to a given player: in this case trust went up an average of 2.6 points $((560 - 266)/111)$, versus only 1.3 for favorable trades $((653 - 511)/111)$. Put another way, if players A and B participate in a trade where player A gives more resources to player B than he or she receives

in return, player A's trust for player B will go up more than twice as much as player B's trust for player A will. What these simple facts show is a correlation between increased trust and participation in a trade, and a stronger such correlation when the trade is imbalanced. This is important because it provides one possible factor which can affect trust in a dynamic environment. It is surprising that a player is more likely to trust another player after participating in an economic transaction that was mathematically unfavorable to them, and this confusing finding certainly presents a course for further study. Determining the direction of causation also provides a course for future research: we will subsequently investigate whether completed trades can be shown to cause an increase in trust, or whether the trust increases stem from some other source, and the trade is simply a byproduct of the higher level of trust.

Beyond the specific results of this study, we see a number of application areas for this game and line of research. First, understanding how trust evolves among people when there are known adversaries within the group can be applied to helping build, and thwart the development of, trust. If the types of interactions that lead to an expectation of trust are known, they can be taken advantage of in a variety of contexts. Those looking to build more successful teams can encourage trust-building interactions. At the same time, those who want to disrupt the success of a group can mimic the behaviors that lead to trust, even when it is undeserved. These teams may be present in gaming environments, virtual and offline collaborations, and in teams of all types. For analysts, knowledge of how trust and distrust form can be used to identify potentially bad actors, or to build strategies for disrupting the collaborations of bad actors in an environment.

VIII. CONCLUSIONS

In this paper, we have described the framework of a generic game which can be used to study how its players determine trust for their opponents, as well as how those trust determinations affect their behavior, and vice versa. We have also described a specific instantiation of this game, and given some anecdotal examples of the data available to be gathered. We have shown that players' trust for each other is affected by economic interactions between them, and that the effect is stronger when the interaction is not equally beneficial to the two participants. Finally, we have described the most common types of social interactions between the players. Trust is an extremely important component in social, economic, and strategic interactions of all types, though a firm understanding of what trust is, how it forms, and how it evolves over the course of a relationship has proved elusive. We hope that this research serves as a step towards an empirical understanding of human notions of trust and trust dynamics, and we intend to continue to

develop further games and other methods of studying this rich and exciting area.

REFERENCES

- [1] S. Marsh, "Formalising trust as a computational concept," Ph.D. dissertation, University of Stirling, Department of Mathematics and Computer Science, 1994.
- [2] R. Falcone and C. Castelfranchi, "Trust dynamics: How trust is influenced by direct experiences and by trust itself," in *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 2004, pp. 740–747.
- [3] D. Melaye and Y. Demazeau, "Bayesian dynamic trust model," *Multi-Agent Systems and Applications IV*, pp. 480–489, 2005.
- [4] E. Fehr and K. Schmidt, "A theory of fairness, competition, and cooperation*," *Quarterly journal of Economics*, vol. 114, no. 3, pp. 817–868, 1999.
- [5] B. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Havlin, "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*. IEEE Computer Society Washington, DC, USA, 2004, pp. 782–789.
- [6] S. Kraus, A. Pfeffer, and S. Shieber, "Colored Trails: A Formalism for Investigating Decision-making in Strategic Environments," *Reasoning, Representation, and Learning in Computer Games*, p. 25, 2005.
- [7] J. Letchford, V. Conitzer, and K. Jain, "An Ethical Game-Theoretic Solution Concept for Two-Player Perfect-Information Games," *Internet and Network Economics*, pp. 696–707, 2008.
- [8] S. Ficici and A. Pfeffer, "Modeling how humans reason about others with partial information," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 315–322.
- [9] Y. Gal, F. Marzo, B. Grosz, and A. Pfeffer, "Learning social preferences in games," in *Proceedings, Nineteenth National Conference on Artificial Intelligence: July, 2004*, pp. 25–29.
- [10] Y. Gal and A. Pfeffer, "Predicting people's bidding behavior in negotiation," in *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. ACM, 2006, p. 376.
- [11] N. Griffiths and M. Luck, "Coalition formation through motivation and trust," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM, 2003, pp. 17–24.
- [12] R. Arunachalam, J. Eriksson, N. Finne, S. Janson, and N. Sadeh, "The TAC supply chain management game," *Computer Science Technical Report CMU-CS-03-184*, Carnegie Mellon University, 2003.
- [13] N. Bos, N. Shami, J. Olson, A. Cheskin, and N. Nan, "In-group/out-group effects in distributed teams: an experimental simulation," in *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. ACM, 2004, pp. 429–436.
- [14] S. Kraus, O. Shehory, and G. Taase, "Coalition formation with uncertain heterogeneous information," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM, 2003, p. 8.
- [15] L. Von Ahn, "Games with a purpose," *Computer*, vol. 39, no. 6, pp. 92–94, 2006.
- [16] L. Von Ahn, R. Liu, and M. Blum, "Peekaboom: a game for locating objects in images," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 2006, p. 64.
- [17] L. Von Ahn, S. Ginosar, M. Kedia, R. Liu, and M. Blum, "Improving accessibility of the web with a computer game," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 2006, p. 82.
- [18] L. Von Ahn, M. Kedia, and M. Blum, "Verbosity: a game for collecting common-sense facts," in *Proceedings of the SIGCHI conference on Human Factors in computing systems*. ACM, 2006, p. 78.

- [19] L. Von Ahn and L. Dabbish, "Labeling images with a computer game," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2004, pp. 319–326.
- [20] E. Law, L. Von Ahn, R. Dannenberg, and M. Crawford, "Tagatune: A game for music and sound annotation," in *International Conference on Music Information Retrieval (ISMIR07)*. Citeseer, 2003, pp. 361–364.
- [21] G. Walsh and J. Golbeck, "Curator: a game with a purpose for collection recommendation," in *Proceedings of the 28th international conference on Human factors in computing systems*. ACM, 2010, pp. 2079–2082.
- [22] Wikipedia, "Mafia (party game)," [http://en.wikipedia.org/wiki/Mafia_\(party_game\)](http://en.wikipedia.org/wiki/Mafia_(party_game)).
- [23] G. Chittaranjan and H. Hung, "Are you awerewolf? detecting deceptive roles and outcomes in a conversational role-playing game," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, 2010, pp. 5334–5337.
- [24] S. Stemler, "An overview of content analysis," *Practical assessment, research & evaluation*, vol. 7, no. 17, pp. 137–146, 2001.