A Model for Competence and Integrity in Variable Payoff Games

Michael J. Smith and Marie desJardins

University of Maryland Baltimore County
1000 Hill Top Circle
Baltimore, MD 21250
{msmith27,mariedj}@csee.umbc.edu

Abstract. Agents often have to trust one another when engaging in joint actions. In many cases, no single design team has the authority to assure that agents cooperate. Trust is required when agents hold potentially different values or conflicting goals. This paper presents a framework and some initial experiments for decomposing agent reputation within a multi-agent society into two characteristics: competence and integrity. The framework models competence as the probability of successfully carrying out an intended action. Integrity is modeled as a rational commitment to maintaining a reputation, based on the agent's assessment of the game's discount rate. We show that a simple, one-level-deep recursive model—given accurate knowledge of self and the other agent's competence and integrity (commitment to reputation)—outperforms tit-for-tat and other standard strategies in evolutionary round-robin iterated prisoner's dilemma tournaments. This indicates that the approach taken here warrants further investigation using more realistic and complex strategies.

1 Introduction

This paper defines a framework for investigating two specific components of trust and reputation, which have not been clearly distinguished in previous decision-making mechanisms:

1. **Integrity:** Is the agent honest about its intentions and commitments?
2. **Competence:** Is the agent capable of fulfilling its stated commitments?

Agents must have some means to decide whether to trust one another when engaging in joint actions. In many cases, no single design team has the authority to assure that agents cooperate. In such environments, dishonest ("defecting") actions may be perceived by some agents to lead to higher utility in the short term. At other times, playing for longer periods of time against opponents with sufficient patience and memory (e.g., extended iterated prisoner's dilemma (IPD) games), the same agent may choose to cooperate. Our goal is to develop models that will allow an agent to decide whether or not to trust another agent, by predicting whether that agent will behave cooperatively in the current situation.

The *competence* of an agent in performing actions is its ability to volitionally change the state of the agents' shared environment—i.e., to correctly carry out its intended actions. *Integrity* is based on a reputation for honesty and a continuing commitment to being trustworthy—i.e., honoring a commitment to take an action. In our framework, an agent can decide whether to act reputably or not, whereas competence is not changeable (or at least not changeable at the time of the decision) by an agent.

A legendary sports story helps to illustrate this distinction. A group of conspirators on the Chicago Black Sox baseball team of 1919 were quite capable of playing well. However, they were not always committed to performing well during the World Series: the conspirators were able to throw games if they believed that it would benefit them. Over the course of that series, they deliberately lost or won, depending on whether the gamblers who were bribing them fulfilled their commitment to pay the stated amount. Because the *competence* level of these players was known to be high, their low performance in some games was quite suspicious to other players on the Black Sox team, who correctly attributed the sporadic losses to a variable *integrity* level—leading to broken trust relationships and reported fights.

Integrity and competence are distinct and should be treated as separate components for decisions based on trust and reputation models. Our experiments show that making a distinction between integrity and competence is beneficial in variable payoff games. Specifically, we show that agents with accurate estimates of the other agents' commitment to reputation (belief in the discount rate for the game) and competence can outperform strategies that do not model these factors.

The ultimate goal of our research is to provide agents with a means to make trusting decisions and a theoretical basis for learning about competence, integrity, and the distinction between the two in multi-agent systems (MAS). An agent is faced with two tasks in using our framework:

1. Observe the behavior of agents in an environment and infer knowledge about their competence and integrity.
2. Apply that knowledge to make effective decisions when interacting with the agents observed.

The experiments in this paper address only the second task of this research agenda: they show the utility of modeling integrity and competence in a simple heterogeneous multi-agent environment. We are currently developing techniques for learning competence and integrity.

2 Prior Work

There is a vast amount of prior work in game theory, probability, and trust modeling that forms the basis of our research. In modeling trust and reputation, there has been a progression from the reciprocal altruism of Trivers [1] to indirect and cyclic reciprocity models. Hamilton identified a key rule for the evaluation of kinship altruism [2]. The image scoring of Nowak and Sigmund is an example of work on societal trust and evolutionary stable systems [3].

Extending the original prisoner's dilemma, Axelrod popularized the study of the iterated prisoner's dilemma (IPD) and the evolution of cooperation [4], and set forth a rationale for the dominance of the tit-for-tat (TFT) strategy. However, with the introduction of noise into the game, TFT can be beaten by another relatively simple strategy, Pavlov [5] [6].

The notion of decomposing trust is not a new one. Marsh treats competence similarly to our decomposition, but does not consider integrity [7]. Instead, he focuses on trust as a subjective probability and adds factors for risk and importance to his formulations. McKnight and Chervany synthesized a high-level typology of trust, based on a broad survey of trust literature; competence and integrity are two of their primary categories of trust, along with benevolence and predictability [8]. Castelfranchi and Falcone note honesty as one of a large number of belief components; their work uses a fuzzy cognitive map model [9]. Schotter notes the relation of risks and incentives for dishonest behavior in an expert vs. consumer game with variable expert competence [10]. In contrast with the above work, our research applies a general framework founded on decision theory to explicitly model and separate competence and integrity and addresses higher incompetence (noise) levels.

Our experiments assume that when our agents use competence or integrity estimates in their decision logic they have perfect knowledge of their own, and other agents', competence and integrity. Clearly this is not a reasonable assumption in real environments. Our ongoing work focuses on developing learning techniques that agents can use to estimate other agents' competence and integrity.

3 Approach

Our trust-based decision framework separately models integrity and competence, combining these factors in the decision-making process. The framework assumes that the environment has the following characteristics:

1. The environment includes a mechanism for identity verification.
2. There are no enforcement or policy mechanisms.
3. Agents will act with bounded rationality.
4. Kinship is not a factor; there is no collusion among agents.
5. There is no oracle for reputation advice.

Agents may model the competence of themselves and/or of other agents, and may model their own and/or other agents' beliefs about how long a given game will last. An agent's competence is modeled as a probability c. Game length is determined stochastically and is modeled as a discount rate γ. A game's γ is simply the probability that another turn will take place after the current iteration. This is equivalent to the w value of Axelrod [4].

In our framework, integrity is this commitment to building a reputation, and is represented in the framework by agents in terms of estimates of γ. We represent an agent's *commitment to reputation* as being directly related to its belief that future encounters will be influenced by the information available from prior encounters with other agents. In general, the lower the perceived game γ—or belief in another player's perception of γ—the more likely the temptation to cheat will overcome any advantage long-term cooperation.

In our framework, competence is represented by agents as a probability estimate, denoted by $\hat{c}$. The competence of the individual agents must be taken into account for each decision and its likely consequences. To simulate simple environments with noise, such as that explored in IPD games [6] [11], where all agents are typically assigned the same error rate in effecting their chosen actions; c is effectively the same for all agents.

Beliefs about agents' competence and integrity could be recursively modeled to an arbitrary depth, as in Gmytrasiewicz and Durfee's recursive modeling of truthfulness [12]. In the experiments described here, we limit our investigation into recursive modeling of competence and integrity to depth one.

The environment we use to investigate the efficacy of our modeling framework is a variation on the iterated prisoner's dilemma (IPD) [4]. Each iteration (two-player game) has variable length. As mentioned earlier, the distribution of game lengths is stochastically controlled by the game discount rate γ. The base payoff matrix of each round of the IPD is shown in Table 1, from the perspective of Player 1 (whose actions are the row labels; Player 2's actions are the column labels). The payoffs are symmetric. This matrix corresponds to the "classic" IPD in relative terms.

Table 1. Prisoner's Dilemma Payoff Matrix

	Cooperate	Defect
Cooperate	R=1	S=-2
Defect	T=3	P=-1

The modification we make to the classic IPD is that the values given in the payoff matrix are multiplied by a *payoff multiplier* M, which is generated randomly for each individual round of a game. Agents know the current multiplier before a committing to an action for each round of play. The payoff multiplier is generated using an exponential distribution: $M_t = -log(1 - random(t))$.

This distribution creates the variation necessary for "confidence game" strategies to be successful, which can cooperate on low-value rounds and then cheat (defect) on high-value rounds to "cash in" on the high payoff. The expected payoff multiplier, $\overline{M}$, is 1.

Each player is assumed to make a commitment to cooperate before each individual round of game play. The game payoff matrix used for the experiments justify the assumption explicitly. A rational player would only choose to play with others who commit to cooperate—since a positive payoff can result only from the opponent cooperating. This is not a limitation for the framework; it is a convention used to simplify the demonstration of the framework's utility. The ultimate role of this trust decision framework is to determine not just what action to take given an interaction, but when to interact with other agents.

3.1 Decision Strategies

We compared three baseline (classic) and three trust-based strategies. The classic strategies are Always-Cooperate (ALL-C), Always-Defect (ALL-D), and Tit-for-Tat (TFT, which matches an opponent's last move, but starts by cooperating). The trust-based strategies model competence and integrity. They either use both agents' estimates of gamma ("Both γ Both c" - BGBC), only their own γ ("Self γ Both c" - SGBC), or only the other agent's γ ("Other γ Both c" - BGBC).

The trust-based strategies tested here always have an estimate of both player's competencies. (In preliminary experiments, we found that *not* modeling competency seldom improved performance, so we have omitted those strategies lacking competence estimates from the experiments reported here.)

What distinguishes the three trust-based strategies from one another is the scope of their knowledge about γ estimates. SGBC agents know only their *self* γ estimate, $\hat{\gamma}_s$. OGBC agents know only the *other* agent's γ estimate, $\hat{\gamma}_o$. BGBC agents have both estimates of their own $\hat{\gamma}_s$ and the other agent's $\hat{\gamma}_o$. In these experiments, the estimates are always correct: the goal is to show that using accurate estimates of competence and integrity can improve performance.

It should be noted that the concept of integrity and its essential nature is not fully captured by this IPD example. If agents were interacting in a more complex environment, competence and integrity would not necessarily be independent.

3.2 The Trust-Based Players' Decision Rule

The SGBC and OGBC players have only one discount rate to apply to their decisions ($\hat{\gamma} = \hat{\gamma}_s$ or $\hat{\gamma}_o$, respectively). When an agent has both estimates, such as a BGBC player, it will compute $\hat{\gamma} = min(\hat{\gamma}_s, \hat{\gamma}_o)$, since either the game's probable length is short and one should defect, or the opponent thinks so, and one should likewise defect.

Each of the trust-based agents uses $\hat{\gamma}$ and $\hat{c}_s$ and $\hat{c}_o$ to compute an estimate of the long-term return of choosing to defect ($\hat{V}_{defection}$) and to cooperate ($\hat{V}_{cooperation}$). The results are then used in the simple decision rule set:

- If $\hat{V}_{defection} > \hat{V}_{cooperation}$, then $defect$.
- If $\hat{V}_{defection} <= \hat{V}_{cooperation}$, then $cooperate$.

The next two sections will describe how $\hat{V}_{defection}$ and $\hat{V}_{cooperation}$ are estimated.

3.3 Discounted Returns

The players compute $\hat{V}_{defection}$ and $\hat{V}_{cooperation}$ as the total estimated return for that action, discounted over the rest of the game, using the competence-influenced estimates of the payoff matrix ($\hat{R}$, $\hat{T}$, and $\hat{P}$—these payoff estimates are described in the next section).

The discounted values of defection and cooperation are given by:

$$\hat{V}_{defection} = \hat{T} + \frac{\hat{\gamma}\hat{P}}{(1-\hat{\gamma})} \tag{1}$$

$$\hat{V}_{cooperation} = \frac{\hat{R}}{1-\hat{\gamma}} \tag{2}$$

The discounted defection estimate given in Equation 1 is pessimistic, a worst case for defection—anticipating a "grim" opponent, who once defected upon will continue to defect for the rest of the game. If the opponent's memory is short (or he discounts old encounters), this overestimates the penalty that an agent will receive for defecting. Equation 2 is optimistic and assumes continuing mutual cooperation if the trust-based player cooperates. Both assumptions are biased towards cooperation. In practice, this effect is mitigated in strategies BGBC and OGBC by their knowledge of γ_o, which should be low if the other player frequently defects intentionally.

In the special case of $\hat{\gamma} = 1$, the assumption is that the game will go on forever, in which case cooperation is the best course of action for the rational agent. (BGBC with good estimates is safe in this case, since it computes $\hat{\gamma} = min(\hat{\gamma}_s, \hat{\gamma}_o)$, but SGBC with $\hat{\gamma} = 1$ is ineffectual.)

The expected payoff multiplier $\overline{M}$ alters Equations 1 and 2 as follows:

$$\hat{V}_{defection} = M_{t_0}\hat{T} + \overline{M}\frac{\hat{\gamma}\hat{P}}{(1-\hat{\gamma})} \tag{3}$$

$$\hat{V}_{cooperation} = M_{t_0}\hat{R} + \hat{\gamma}\overline{M}\frac{\hat{R}}{1-\hat{\gamma}} \tag{4}$$

3.4 Estimated Payoffs

The estimate of payoffs for the four possible outcomes in the payoff matrix must be adjusted by the self-estimate of an agent's own competence ($\hat{c}_s$) and its estimate of the other agent's competence ($\hat{c}_o$). Equation 5 gives the expected payoff if both agents cooperate, $\hat{R}$. Since the agents have imperfect competence, any of the four outcomes is possible. Therefore, the expected payoff is the sum of the four outcomes, weighted by their respective probabilities. For a given joint intent (R, S, T, or P), which corresponds to the intersection of the two players' individual intents in Table 1, the estimated competencies of the agents ($\hat{c}_s$ and $\hat{c}_o$) determine the probability of each actual outcome. For example, if the joint intent of the agents is R (both cooperate), then the temptation payoff T (i.e., I defect despite intending to cooperate, but you cooperate as intended) occurs with probability $(1-\hat{c}_s)\hat{c}_o$. The expected payoffs then are:

- Estimated payoff for case where both cooperate

$$\hat{R} = \hat{c}_s\hat{c}_o R + \hat{c}_s(1-\hat{c}_o)S + (1-\hat{c}_s)\hat{c}_o T + (1-\hat{c}_s)(1-\hat{c}_o)P \tag{5}$$

– Estimated payoff for defecting on a cooperating opponent

$$\hat{T} = \hat{c}_s \hat{c}_o T + \hat{c}_s(1 - \hat{c}_o)P + (1 - \hat{c}_s)\hat{c}_o R + (1 - \hat{c}_s)(1 - \hat{c}_o)S \quad (6)$$

– Estimated payoff for case where both defect

$$\hat{P} = \hat{c}_s \hat{c}_o P + \hat{c}_s(1 - \hat{c}_o)T + (1 - \hat{c}_s)\hat{c}_o S + (1 - \hat{c}_s)(1 - \hat{c}_o)R \quad (7)$$

Note that for the final computation of the discounted payoff after cooperation and the discounted payoff after defection, $\hat{S}$ is never needed in the decision rule that follows, since nobody *expects* to be a sucker. However, the computation of $\hat{S}$ would be analogous.

4 Experiment Design

We used an evolutionary computing model to evaluate the strategies in a long-term competitive environment. This ensured we could establish the stability of the strategies in dynamically changing populations. As some of our outcomes show, an initially effective strategy can be overcome by other strategies over time as other strategies wane (the evolutionary model shown here also lends itself to exploring evolutionary learning strategies in future work.) Agents were ranked individually, and new populations of the strategies were selected in each generation with linear ranking [13] and stochastic universal sampling [14]. The fitness function is simply the average payoff for the games played on the most recent round. In this process, the selection bias b is a control on the rate of change. An individual i of n players is ranked r_i by its fitness score and is selected with probability $p_i = \frac{b}{n} - (r_i - 1)\frac{2(b-1)}{n(n-1)}$. Based on preliminary experiments, we chose $b = 1.1$, which provided more reliable results within a reasonable time. At the end of each generation there was total replacement of individuals.

An exponential multiplier was applied to the payoffs for each round's decision, as described previously. Both players knew the full payoff matrix and were informed of the current multiplier before selecting an action.

The agents' knowledge of γ values and c values were always accurate for the trust-based strategies. For the simulations, perfect knowledge of TFT and ALL-C was modeled by using an estimate of $\gamma = 1$ and ALL-D was estimated as having $\gamma = 0$. It can be argued that TFT is not perfectly modeled by $\gamma = 1$, since it does sometimes defect. However, given a cooperating opponent, TFT instantaneously reverts to $\gamma = 1$: it will never defect first, no matter how high the payoff, and thus behaves as if it had $\hat{\gamma} = 1$. Note, however, that a TFT agent with c lower than 1 *will* sometimes defect; modeling γ and c separately allows agents to differentiate unintended and intended defection.

In all of the experiments, the initial populations consist of 50 individuals per strategy. Each game was run for 1000 generations. In each generation, every player played one IPD game against every other player in round-robin fashion. Each IPD game had a random number of rounds, determined

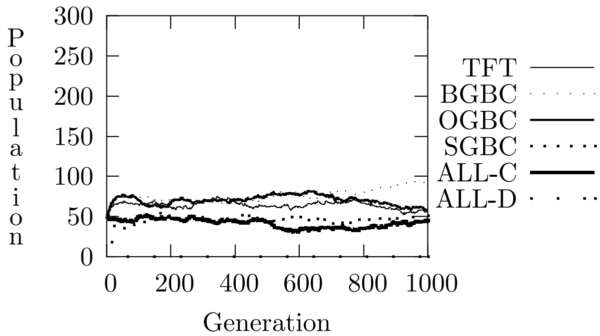

Fig. 1. Competition with $b = 1.1$, $\gamma = \hat{\gamma} = 0.9$, $c = \hat{c} = 1.0$. High γ inhibits decision rule of trust-based strategies.

stochastically by the game's true γ as described earlier. Twenty-five experiments were conducted, using the Cartesian product of true discount factors γ selected from $\{0.95, 0.9, 0.8, 0.7, 0.6\}$ and agent competence c selected from $\{1.0, 0.9, 0.8, 0.7, 0.6\}$. In these experiments, all agents have the same competence, making them comparable to typical IPD experiments with noise.

These ranges were sufficient to identify key phase transitions between successful strategies. Note that a c value of less than 0.5 does not make sense for the 2x2 game payoff matrix, since an agent with very low competence could simply attempt to do the opposite of the desired result (and be more successful than another agent with slightly higher confidence).

4.1 Experiments

In the case with high discount rate and high competence ($\gamma = 0.9, c = 1.0$, Figure 1), there is no compelling advantage to any strategy, except that ALL-D consistently loses, becoming extinct within a few generations. By contrast, with high c but a lower $\gamma = 0.8$ (Figure 2), the strategies that model at least the agent's own γ quickly eliminate all of the other strategies. Interestingly, the OGBC strategy, which models the other agent's γ but not its own, performs poorly. TFT, ALL-C and ALL-D all die out.

The high-competence (effectively noise-free) environment lends itself to the trust-based strategies quite well. As γ decreases, the trust-based strategies outperform the standard strategies, including TFT. However, when we decrease competence, the trust-based strategies perform less well. In Figure 3, we see that keeping $\gamma = 0.8$ but lowering c to 0.8 results in a resounding success for TFT, which extinguishes most of the strategies, although BGBC and SGBC do manage to survive in low numbers. Figure 4 shows the transition from dominance by TFT to dominance of the trust-based strategies (again, specifically BGBC and SGBC) as γ is lowered again, from 0.8 to 0.7.

The answer to TFT's success when competence is low lies in a weakness in the decision framework logic: specifically, the estimate of the payoff for the

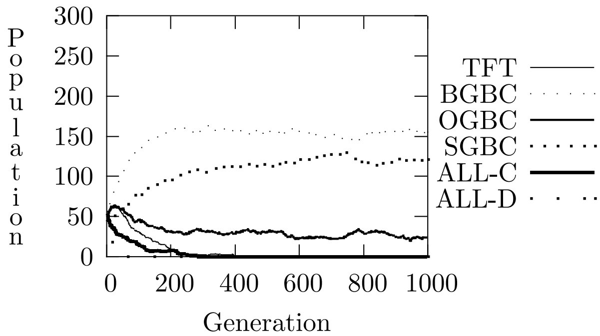

Fig. 2. Competition with $b = 1.1$, $\gamma = \hat{\gamma} = 0.8$, $c = \hat{c} = 1.0$. Trust-based strategies eliminate other strategies with decreased γ; decision rule effective

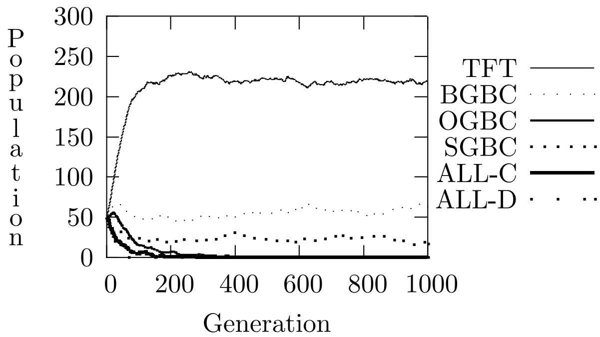

Fig. 3. Competition with $b = 1.1$, $\gamma = \hat{\gamma} = 0.8$, $c = \hat{c} = 0.8$. TFT most fit, but BGBC and SGBC survive

discounted defection. Equation 3 is overly pessimistic in its second term, which takes the worst case, assuming that any player who is defected upon will become "grim" and continuously defect for the remainder of play. In fact, TFT and many other strategies have components of forgiveness (or forgetfulness), and forgiveness is a frequently commented on as an aspect of trust and reputation theories [4] [15]. However, this factor only explains why the trust-based strategies don't *outperform* TFT, not why TFT beats them so soundly.

A second weakness explains TFT's advantage here. TFT's retaliatory nature has no counterpart in the trust-based strategies. TFT always responds with a "tat" defection, even after unintentional defections. This usually results in the sucker's payoff S.

In addition, the few high-payoff cases a trust-based player can exploit now have uncertain execution—which would otherwise make the first term of Equation 3 outweigh the pessimistic second term, causing the trust-based players to defect—so TFT tends to overtake the trust-based strategies in the individual rankings.

Another interesting case is that of with high γ and moderate c; in this scenario, the classic strategies beat the trust-based strategies, and must compete

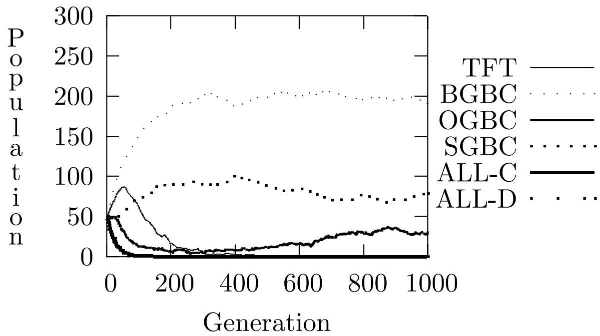

Fig. 4. Competition with $b = 1.1$, $\gamma = \hat{\gamma} = 0.7$, $c = \hat{c} = 0.8$. BGBC, SGBC, and OGBC drive out other strategies

against each other. With γ in the range from 0.95 to 0.9, either the classic strategy TFT or ALL-D will win. TFT wins when the trust-based strategies cooperate enough to allow TFT to outscore ALL-D. If they are eliminated along with ALL-C, TFT loses, as ALL-D gets the reward for defection the first play.

As γ lowers, the trust-based strategies become effective in the presence of significant incompetence (noise). The trend of effectiveness for the trust-based strategies continues as γ decreases—despite the imperfect assumptions of their decision rule—and they begin to dominate TFT. When γ drops towards 0.6, BGBC dominates the other players by exploiting them frequently. SGBC and OGBC may survive above $\gamma = 0.6$, but are eliminated when the time horizon is so short.

Taken as a whole, the 25 experiments demonstrate the potential of our decompositional approach to trust modeling, and highlight areas that require further refinement in the framework and in the decision logic. Where effective, the trust-based strategy with the most information, BGBC, does best, followed by SGBC, and then OGBC. SGBC outperforms OGBC because SGBC has the true γ, whereas $\hat{\gamma}_o$ of the simple classic strategies is of comparatively little value. In general, one finds that as γ increases in the presence of error, the trust-based strategies perform less well, and are eventually outperformed, usually by TFT. This phase transition occurs because of the underestimate in second term of Equation 3 for any strategy (such as TFT) that is forgiving and shows an absence of retaliatory behavior.

5 Future Work

We plan to move away from deterministic strategies and evaluate more general agents that have bounded rationality. The decision-theoretic framework will be adapted to determine whether an agent wants to interact with another player, which addresses the non-protocol aspects of Ramchurn's desiderata [16]. Some additional classic strategies may be considered, especially variations of Pavlov, which is known to perform well in noisy environments.

A key area for future work is addressing some form of discounting associated with forgiveness and forgetfulness in the framework. This would mean experiments with morality- or memory-based discounting, respectively. These discounts would be modeled in addition to the $\hat{\gamma}$ discount accorded to the expectation of there being a future game.

Additional experiments may be conducted to further establish a firm empirical basis for the framework's viability when agents possess imperfect estimates of c and γ. Since observing other agents and arriving at accurate beliefs for γ_s, γ_o, and c_o will be difficult, the knowledge of those parameters will presumably be far less than perfect.

We plan to explore the theoretical implications of the recursive modeling of agents' beliefs about each other. We believe that it may be possible to show that one can mitigate or even eliminate the necessity of recursive modeling for integrity and competence beliefs for some classes of MAS applications. It will also be helpful to exploit notions of contextual trust and ontological reasoning within a network of interacting agents may help to assist in determining $\hat{\gamma}$ and $\hat{c}$ through abstraction and generalization.

Our ultimate goal is to demonstrate an effective means to learn integrity and competence estimates and to deconflict the two in a simulation setting. Discovering methods for an agent in a MAS to learn and then apply estimates of other agents' integrity and competence appears to be possible under at least some conditions. A determination of competence at a particular skill could be learned by observing attempts to use a skill where it would be surely be advantageous to an agent to succeed. With enough variation in the payoff matrices, integrity and competence may be learned as separate parameters by observing other agents' interactions in this environment.

The work of Mui [17], Whitby [18], Sabater and Sierra [15], and Yu and Singh [19] in reputation propagation, evidence, and learning will provide elements for our current research on learning estimates of other agents' competence and integrity, as well as other parameters we may wish to model.

6 Conclusion

The motivation for well grounded models of trust and reputation that can be applied to a wide variety of MAS is increasing as multi-agent environments become more common and larger-scale. Trust and reputation problems modeled in MAS involving joint intention and action have real-world, commercial analogs including e-commerce market reputation systems, contracting, and supply-chain management.

The experiments presented in this paper demonstrated that an explicit framework based on a fairly straightforward application of decision theory principles, combined with a framework for separating competence from integrity (commitment to reputation), can be effective over much of the space of possible competences and game discount rates. Simple extensions to the framework to incorporate some sense of the forgiveness (or forgetfulness) of other players should eliminate a key weakness identified in the decision rule provided as part of the framework here.

References

1. Trivers, R.: The evolution of reciprocal altruism. Quarterly Review of Biology **46** (1971) 35–57
2. Hamilton, W.: The genetical evolution of social behaviour. i. Journal of Theoretical Biology **7** (1964) 1–16
3. Nowak, M., Sigmund, K.: Evolution of indirect reciprocity by image scoring. Nature **393** (1998) 573–577
4. Axelrod, R.: The Evolution of Cooperation. Basic Books, New York (1984)
5. Kraines, D., Kraines, V.: Pavlov and the prisoner's dilemma. Theory and Decision **26** (1989) 47–79
6. Kraines, D., Kraines, V.: Learning to cooperate with pavlov: An adaptive strategy for the iterated prisoner's dilemma with noise. Theory and Decision **35** (1993) 107–150
7. Marsh, S.P.: Formalising Trust as a Computational Concept. PhD thesis, University of Stirling (1994)
8. McKnight, D.H., Chervany, N.L.: Trust and distrust definitions: One bite at a time. In: Trust in Cyber-societites, Integrating the Human and Artificial Perspectives, London, UK, Springer-Verlag (2000) 27–54
9. Castelfranchi, C., Falcone, R., Pezzulo, G.: Trust in information sources as a source for trust: A fuzzy approach. In: Proceedings of AAMAS '03, Melbourne, Australia, ACM (2003) 89–96
10. Schotter, A.: On the Economic Virtues of Incompetency and Dishonesty. In: Paradoxical Effects of Social Behavior. Springer-Verlag (1986) 235–241
11. Wu, J., Axelrod, R.: How to cope with noise in the iterated prisoner's dilemma. Journal of Conflict Resolution **39** (1995) 183–189
12. Gmytrasiewicz, P., Durfee, E.: Toward a theory of honesty and trust among communicating autonomous agents. Group Decision and Negotiation **2** (1993) 237–258
13. Branke, J.: Evolutionary Optimization in Dynamic Environments. Kluwer Academic Publishers (2002)
14. Baker, J.E.: Reducing bias and inefficiency in the selection algorithm. In: Proc. of the 2nd Intl Conf on GA, Lawrence Erlbaum Associates, Inc. Mahwah, NJ, USA (1987) 14–21
15. Sabater, J., Sierra, C.: Social regret, a reputation model based on social relations. In: ACM SIGecom Exchanges. Volume 3.1. ACM (2002) 44–56
16. Ramchurn, S.D., Hunyh, D., Jennings, N.R.: Trust in multi-agent systems. Knowledge Engineering Review (2004)
17. Mui, L., Mohtashemi, M., Ang, C., Szolovits, P., Halberstadt, A.: Ratings in distributed systems: A bayesian approach. In: Workshop on Information Technologies and Systems. (2001)
18. Whitby, A., Josang, A., Indulska, J.: Filtering out unfair ratings in bayesian reputation systems. In: Proceedings of the 7th Int Workshop on Trust in Agent Societies. (2004)
19. Yu, B., Singh, M.P.: An evidential model of distributed reputation management. In: Proceedings of First International Conference on Autonomous Agents and MAS, Bologna, Italy (2002) 294–301