Effects of Belief and Memory on Strategic Negotiation

Markus Guhe
Alex Lascarides
School of Informatics
University of Edinburgh
Edinburgh EH8 9AB, Scotland
m.guhe@ed.ac.uk
alex@inf.ed.a.uk

Kevin O'Connor
Verena Rieser
School of Mathematical and Computer Sciences
Heriot Watt University
Edinburgh EH14 4AS, Scotland
kho30@hw.ac.uk
v.t.rieser@hw.ac.uk

Abstract

We present an empirical framework for testing negotiation strategies in a complex win–lose game that lacks any analytic solution. We explore how different belief and memory models affect trading and win rates. We show that cognitive limitations can be compensated for by being an 'optimistic' negotiator: make your desired trade offer, regardless of your beliefs about how opponents will react. In contrast, agents with good cognitive abilities can win with fewer but more effective offers. Corpus analysis shows human negotiators are somewhere in between, suggesting that they compensate for deficient memory and belief when necessary.

1 Introduction

Strategic negotiation is a type of non-cooperative conversation, which the Gricean view of cognitive agents fails to account for (Asher and Lascarides, 2013). In this paper we investigate negotiation dialogues as they occur during trading, within the example domain of the board game *The Settlers of Catan* (or *Settlers*, Teuber, 1995; see `www.catan.com`. We explore how human errors in beliefs, in particular forgetting, impact negotiating and trading behaviour, and take first steps towards building negotiation strategies that are effective in spite of deficiencies in beliefs.

Trading and bargaining is often modelled as rational actions between agents, all of whom maximise their *expected utilities*—an optimal trade-off between what they *prefer* (typically defined by a utility function) and what they *believe* they can achieve (typically defined via a dynamic Bayesian network; Savage, 1954). Solving a game problem involves finding *equilibrium strategies*: an optimal action for each player in that it maximises his expected utility, assuming that the other players perform their specified action (Shoham and Leyton-Brown, 2009). But this Savagean model of decision making is highly idealised and humans often deviate from it (Kahneman and Tversky, 1979; Ariely, 2008; Yong and Xinlin, 2012). Non-optimal human behaviour occurs in complex games, for instance, where existing algorithms for computing expected utilities can also break down. *Settlers* is one such game. It involves players conversing to negotiate trades over restricted resources. And even if an analytic solution for trading in *Settlers* were to exist, it doesn't necessarily match what humans do.

One response to this is to develop a symbolic model consisting of *heuristics* that match the strategies of expert human players (Thomas, 2003). But their effectiveness and correlation to human behaviour must be evaluated. Accordingly, we present here an empirical framework for devising and evaluating heuristics. We focus on testing agents with various cognitive limitations (e.g., memory loss), and we show that limited cognitive abilities can be compensated for by being an optimistic negotiator: make an offer for your most desired trade whatever your beliefs about its outcome. We then compare various computational agents with a corpus of humans playing *Settlers* (Afantenos et al., 2012).

2 The Settlers of Catan

Settlers is a win–lose board game for 2 to 4 players. Each player acquires resources (ore, wood, wheat, clay, sheep) and uses them in different combinations to build roads, settlements and cities on a board like the one shown in Figure 1. This earns points and the first player with 10 points wins. Players acquire resources in several ways, e.g., via the dice roll that starts each turn and through trading with other players—so players converse to negotiate trades. The dice rolls make

Proceedings of the 17th Workshop on the Semantics and Pragmatics of Dialogue, December 16-18, 2013, Amsterdam, The Netherlands.

Figure 1: A game of *Settlers* in *JSettlers*.

future states non-deterministic, compelling players to assess the risk of their moves, including trading moves. A player's decisions about what resources to trade depends on what he wants to build; e.g., a road requires 1 clay and 1 wood. Trading decisions are also determined by estimates of what will most advance, or undermine, the opponents' strategies (Thomas, 2003). Players can also lose resources: e.g., a player who rolls a 7 can rob from another player. What's robbed is hidden from view, so players lack complete information about their opponents' resources. Because *Settlers* is a game of imperfect information, agents can, and frequently do, engage in 'futile' negotiations, which don't result in any trade. An agent that initiates a negotiation that doesn't result in a trade has in effect miscalculated the equilibrium strategies.

3 Related Work

There are several empirical approaches to modelling *Settlers*, but none of them includes trading or negotiation. Szita et al. (2010) and Roelofs (2012) use Monte Carlo Tree Search (MCTS), but on a simplified version of *Settlers* without any negotiation and trading between players. In contrast, our focus is on negotiation strategies and their interaction with cognitive ability.

Pfeiffer (2003) uses reinforcement learning to acquire *Settlers* strategies from game simulations. The results show that a mixture of hand-coded heuristics and learnt strategies challenges human players. This is encouraging, but it is unclear whether the heuristics are cognitively plausible or yield successful trading strategies. We address this here by evaluating different trading heuristics in a simulated game environment and comparing the results with what people do.

Related work on imperfect information games, such as poker and *Settlers*, emphasises the importance of modelling beliefs (Sweeney, 2012). An agent's beliefs about his opponents' intentions is known as *opponent modelling* or *nested beliefs* (Rieser et al., 2012). Vogel et al. (2013) have shown that nested beliefs are useful for reasoning about implicature-rich interpretations in dialogue. So the following experiments evaluate how the capacity to accurately model nested beliefs contributes to negotiating, trading and winning.

4 Planning in *JSettlers*

We use an open source implementation called *JSettlers* (`jsettlers2.sourceforge.net`, Thomas, 2003). *JSettlers* is a client–server system: a server maintains the game state and passes messages between each of the players' clients, which can run on different computers. Clients can be human players or computer agents. Here, we report on simulations between computer agents.

The *JSettlers* agent, which we call the *original* agent, goes through multiple phases after the dice roll that starts its turn:

1. Deal with game events: e.g. placing the robber; acquiring or discarding resources.
2. Determine legal and potential places to build.
3. Estimate the time required to build pieces on legal places (the ETB).
4. Compute the *Best Build Plan* (BBP): a sequence of build actions that achieves 10 points in the shortest estimated time (ignoring how opponents might hinder your plans).
5. Try to execute the BBP, including negotiating and trading with other players and/or the bank or a port.

As we are exploring how various cognitive limitations impact decisions to negotiate and how that affects trading and winning, all our agents adopt the same *build* strategy: i.e., steps 1–4 remain constant, while step 5 differs across agents. We first describe these common steps.

Agents sort resources into *needed* vs. *not needed* given their BBP. When considering whether to offer an unneeded resource for a needed one (and likewise when considering whether to accept an offer addressed to them), an agent compares the *Estimated Build Time* (ETB) of the offer against that of its *Best Alternative to Negotiated Agreement* (BATNA), the latter being no trade at all, a bank trade (at a 4:1 ratio) or a port trade (at a 3:1 or 2:1 ratio, depending on the port). All our agents make the offer only if the former ETB is shorter than the latter.

Sufficient conditions for making an offer differ among our agents, however. By using the same ETBs and build policies, all our agents have the same 'instrinsic' trading preferences, but differ in extrinsic preferences because of different beliefs about the offer's outcome. Their differing beliefs stem from their different *cognitive capacities* (how much evidence they have for predicting an offer's outcome), and their different ways of handling missing evidence (see Section 5 for details).

Agents have three possible responses to a trade offer during a negotiation: to *accept* it (and enact the trade), to *counteroffer* (which may be a completion of a partial offer) or to *reject* it. An agent accepts an offer if the offer is executable and its ETB is less than that of its BATNA; he counteroffers if the offer's ETB is not less than that of its BATNA but there is an offer that he hasn't already made that satisfies the agent's sufficient conditions for offering; otherwise he rejects the offer.

5 Negotiation Strategies

To win, an agent needs an effective *trading* strategy and an effective *negotiation* strategy. An effective trading strategy is one that increases the player's likelihood of winning: i.e., on average, his choices of when and what to trade help him more than they hurt him. An effective negotiation strategy increases the likelihood that negotiations culminate in an effective trade. Our initial experiments demonstrate that the implemented *JSettlers* trading strategy is effective (see Sections 6.2 and 6.3). So in order to evaluate distinct negotiation policies and belief models with appropriate controls, we make all agents adopt the same effective *JSettlers* trading choices.

A trade offer has a non-deterministic outcome because it depends on the opponents reaction: a desired outcome is that the offer is accepted (so the trade is enacted) but without that trade helping the opponent more than it helps the proposer. Here, this means that the opponent has fewer than 8 points and his BBP doesnt block the proposers own BBP. Thus agents should use evidence, both past and present, to estimate the opponents (hidden) resources and BBP, to infer whether an offer will have this desired outcome.

In this paper, we manipulate the cognitive capacity of an agent via how much evidence he has for predicting an outcome; and we manipulate how optimistic or pessimistic he is about an outcome when evidence is inconclusive. So overall, we manipulate what evidence an agent uses, and how he uses it. In total we investigate 10 conditions, see table 1.

On **cognitive capacity**, we investigate at one extreme agents who are *omniscient* about the opponents' resources and/or BBPs (making the game state fully observable), and at the other extreme agents who lack any evidence at all (either past or present) for inferring them. Within these two extremes, we implement agents who use past and present evidence to estimate the opponents' resources and BBPs, but they *forget* past evidence after a certain time. We give the details in Section 6.

We then distinguish three ways of using evidence to predict the outcomes of dialogue moves. An **optimistic proposer** makes a trade offer with the best possible ETB regardless of the evidence for its outcome; thus, he in effect ignores evidence that's against the desired outcome even when it exists!

Alternatively one can be less optimistic, making a trade offer only if the available evidence yields a belief in its desired outcome. In fact, we implement two non-optimistic agents, which handle *missing evidence* differently. The original *JSettlers* agent is what we call a **cautious proposer**: he makes an offer only if he believes it will have the desired outcome, but he defaults to this belief when the evidence is inconclusive. That is, when he has insufficient evidence to infer certain information about his opponents' resources and BBPs, he simply assumes favourable values.

This contrasts with a **pessimistic proposer** who makes a trade offer only if he believes it will have the desired outcome, but unlike the cautious proposer he defaults to believing it won't in the absence of information to the contrary. Thus, by default the pessimist assumes the trading partner does not have his desired resource or is not willing to sell it or does not need what the pessimist is offering.

Finally, we implement agents who issue *partial* offers like (1a), as opposed to only making *complete* offers like (1b):

(1) a. I need clay.
 b. Max, will you give me 1 clay for 1 wood?

The original *JSettlers* agent only makes *complete* offers like (1b), but this isn't human-like (see Section 6.8 for details): only 4.7% of the offers in the corpus we collected of people playing *Settlers* are as specific as (1b); 23.1% specify the resources but not their quantity; and the most frequent type of offer, at 34.8%, is a partial offer specifying the receivable resource, as in (1a).

In addition to partial offers, we evaluate different proposer-types (optimistic, cautious, pessimistic) with different cognitive abilities (omniscient, ignorant, forgetful) in terms of their effectiveness as negotiators and their chances of winning. We started with Thomas' (2003) **original** *JSettlers* agent. We modified the code slightly, e.g., by improving the initial placement of pieces and fixing a number of bugs that, for example, affected the automated running of large numbers of games in our simulations. However, the main planning and trading mechanisms are unchanged, and this agent remains a cautious proposer.

This original agent is neither omniscient nor ignorant, but he's not human-like: He has a perfect memory, never forgetting past evidence, and so maintains a perfect model of the opponents' resources until a 7 is rolled. At this point, the robbed player loses a resource to the robbing player, and any player with 7 or more resources must discard half of them. Which resources a player loses is unobservable to agents not involved in the transfer, and so the *JSettlers* agent downdates his beliefs for *all* resources of the affected player to *unknown*. This extreme form of belief change is also not human-like, as humans would still keep an hypothesis of potential resources owned by this player. Here, we investigate how human memory, in particular forgetting, influences the effectiveness of various distinct negotiation strategies, showing that some of these penalise a player with deficient beliefs while others do not. We thus make the first steps towards building negotiation strategies that are effective in spite of human-like errors in beliefs. In future work, we will also investigate how more elaborate forms of belief update and revision (after unobservable events) will influence negotiation strategies.

The original agent is not strong but is at least in the ballpark of human performance. Thomas (2003) performed an evaluation where in each game three agents (agents without our bug fixes and improvements) played one human player. The human player won about 50% of the games (and each agent about 17%).

6 Experiments with Modified Agents

6.1 Method

Simulations for testing a particular belief model and strategy for proposing trades all consist of 1 modified agent playing 3 original agents in 10,000 games. So the null hypothesis is that each agent wins 25% of these 10,000 games. To carry out these simulations, we created a simulation environment for *JSettlers*. The server and the 4 agents all ran on the same machine, and a simulation of 10,000 games took about 0.5–1h on a desktop computer.

In addition to measuring the win rate, we analyse the agents' negotiating and trading behaviour: the number of offers they made, the total number of successful offers (i.e., how many trade offers resulted in a trade), and the total number of trades with other players, i.e. including trades resulting from accepting other players' offers. Finally, the proportion of an agent's offers that are successful gives a rough measure of how accurately he estimates an equilibrium trading move. Table 1 gives an overview of the results.

Due to the large number of games in each simulation even small differences can be significant (as long as the standard deviation is also reduced). At the same time, in the simulations reported here there are no significant differences between the three instances of the original agent, i.e. all differences result from agent modifications. For the simulations, we test significance of win rates against the null hypothesis (25%) by using the z-test; we analyse differences in trading behaviour among opponents via paired t-tests for all combinations of opponents (in fact, there were no significant differences for the offering/trading measures between any two original opponents); and we use a significance threshold of $p < 0.01$. We now describe the simulations in detail, in the order in which they are given in Table 1.

	modified agent					averages for the 3 original agents				
	win rate	offers	successful offers	total trades	successful offers/ offers	win rate	offers	successful offers	total trades	successful offers/ offers
random BBP	**0.022**	**21.9**	**5.1**	**9.0**	**0.23**	0.326	14.2	3.9	8.1	0.27
non trading	**0.127**	**0.0**	**0.0**	**0.0**	---	0.291	17.7	3.1	6.2	0.18
omniscient · resource	0.244	**7.5**	3.7	7.5	**0.50**	0.252	13.6	3.8	7.6	0.28
omniscient · BBP	0.243	13.8	**3.6**	7.6	**0.26**	0.252	13.7	3.8	7.6	0.28
omniscient · BBP/resource	0.243	**7.3**	**3.5**	7.5	**0.48**	0.252	13.8	3.9	7.5	0.28
ignorant · cautious	0.241	**28.5**	**3.6**	7.7	**0.13**	0.253	13.9	3.9	7.7	0.28
ignorant · pessimist	**0.162**	**0.0**	**0.0**	**4.2**	---	0.279	14.7	4.2	7.0	0.28
ignorant · optimist	0.255	**27.1**	**3.7**	**7.8**	**0.14**	0.249	13.8	3.9	7.6	0.28
ignorant · caut. 65% cap	**0.203**	**12.7**	**1.9**	**6.1**	**0.15**	0.266	14.2	4.0	7.3	0.29
partialising initially	0.240	**19.4**	**2.7**	**7.9**	**0.14**	0.253	14.7	4.2	7.7	0.29

Table 1: Results of 1 modified agent playing 3 original agents. Results for the modified agents are shown in the left half; the corresponding averages of the results for the 3 original agents are shown in the right half. Significant results ($p < 0.01$) are in bold.

6.2 Quality of Build Strategy

An agent's trading and negotiation strategies are based on a best build plan (see step 4 in the agent's planning procedure, Section 4). In order to establish whether we can retain the original agent's build plan for our experiments, we evaluate its quality by testing it against a **random BBP agent**, who chooses his 'best' build plan randomly.

Results and Discussion. This agent wins only 2.2% of the games, even though he makes more trade offers, more successful offers and more trades (using the same negotiation and trade strategy as the original agent). So the original agent's build strategy is an improvement over a random baseline. Together with the evaluation in Thomas (2003), we can therefore assume that all of our agents, which retain the *JSettlers* build strategy, have decent build plans.

6.3 Benefits of Trading

To establish that trading contributes to winning, we created the **non-trading agent**, which is like the original agent except that he never trades with other players but only with the bank or a port.

Results and Discussion. The non-trading agent wins only half as many games as his opponents, providing strong evidence that the *JSettlers* trading policy is effective and contributes to winning. (The agent makes more trades with the bank and ports (9.6) than his original opponents (4.9), but this does not compensate for not trading with agents.) So the agent's *preferences* over possible trades, defined by his ETBs and BBP, correlates with his chances to win. Since we never change these calculations, any changes to win rates will stem from how effective the *negotiation strategy* is in achieving a trade.

6.4 Beliefs: Omniscience

To explore how useful accurate beliefs about the opponents' resources are, we tested a **resource omniscient agent**, who *directly observes* his opponents' resources but remains a **cautious proposer**. So he never has defeasible beliefs about resources, but may default to a belief that his opponent has a favourable BBP. We make the resources observable by getting the (original) opponents to declare them at the start of each turn (note that we don't allow deception in our simulations; the role of deception in *Settlers* is future work). Moves that declare resources, or lack of them, are attested frequently in the human *Settlers* corpus (Afantenos et al., 2012), generally via responses to questions about offers; e.g., *I've got clay* in response to *What will you give me?*, or *I don't have any* in response to *I need wood.*

We also tested a **BBP-omniscient agent**: again a **cautious proposer** but one for whom his opponents' BBPs are always observable. Again, these are observable because agents declare them (*I intend to build a road.*). These are attested but rare in the *Settlers* corpus.

Finally, we test an agent that is **omniscient on both BBPs and resources**.

Results and Discussion. None of these three agents have significantly different win rates than their original opponents. However, they all have a more effective negotiating strategy, with a higher proportion of their offers being successful. But the resource-omniscient agents make far fewer offers overall, and the BBP-omniscient agent makes slightly but significantly fewer successful trade offers.

With respect to human games, entering fewer 'futile' negotiations is relevant because human players can easily get annoyed when players make many trade offers, and in particular offers where it is obvious that it won't be accepted.

It may seem counterintuitive that making hidden parts of the *Settlers* game state observable fails to improve the win rate. We believe that this happens for three reasons. First, the original agent's perfect memory gives him good resource-tracking capabilities: the only relative advantage of the resource-omniscient agent comes after a 7 has been rolled.

Secondly, and perhaps more importantly, all of our agents, both modified and original, are only willing to negotiate for their best possible trade, and do not consider entering a negotiation for a 'second best' possibility should they believe that their best possible trade isn't achievable. This aspect of the negotiation strategy hurts the omniscient agent: since he never defaults to favourable values, he starts a negotiation less frequently than an original agent would (in our simulations, around half the time). This denies the omniscient agents the chance to consider *counteroffers* that may not be best but are nevertheless effective and so acceptable. But the non-omniscient agent gets relatively more opportunities to consider such counteroffers. In future work, we plan to investigate how adapting the negotiation strategy to allow initiating a negotiation for a 'second best' build plan would enhance the win rates.

Finally, making declarations of BBPs seems to be largely redundant: all agents use the same evaluation function for computing everyone's BBPs, and for the opponents' BBPs this function draws only on the observable part of the game state and beliefs about the opponents' (hidden) resources, which even for the original agent are relatively accurate. This redundancy also explains why the **BBP/resource-omniscient** agent performs to much the same level as the resource omniscient agent. In future work, we plan to investigate how declaring build plans when agents have distinct build policies—as humans players invariably do—impacts game performance.

6.5 Beliefs: Ignorant

The **ignorant agent** does not track the other agents' resources at all. He can deal with this lack of knowledge in 3 ways: being **optimistic** (assume all opponents have all resources), being **cautious** (treat all opponents' resources as unknown but track the overall amount, and default to assuming they have the desired resource if this doesn't conflict with knowledge of how many resources they have) or being **pessimistic** (assume all opponents have no resources). So the pessimistic agent never makes an initial offer (because he believes that the offer will be rejected) but may accept offers from others. In contrast, the optimistic and cautious agents make many offers.

Results and Discussion. The ignorant pessimist wins significantly fewer games and makes fewer trades (and no trade offers, by design). In contrast, the cautious and optimistic proposers can compensate for the deficient belief model by making twice as many trade offers. Thus, while it pays to have a relatively accurate belief model and to reason about likely outcomes of offers, there are also negotiation strategies that compensate for a deficient belief model, which involve *ignoring* the risk of your desired offer having an undesired outcome, i.e. not being accepted.

Capping the number of trade offers. Because the optimistic and cautious proposers compensate for ignorance by making many trade offers, we tested whether the driving factor is the quantity or the quality of the offers, by capping the number of offers the cautious ignorant agent can make to the same level as his original opponents. We implemented this by letting the ignorant agent make the decision on whether to make a trade offer but then only 'allowed' it to actually go through with making it 65% of the time. As can be seen in Table 1, this agent fares much worse than the equivalent 'non-capped' ignorant agent, showing that he is able to approach the winning rates of his less belief-deficient opponents only by counteracting the lower quality in his trade offers by increasing their quantity. Note also that in this simulation, the original opponents make fewer trades because the ignorant capped agent makes fewer acceptable offers.

6.6 Negotiation Strategy: Partial Offers

As mentioned before, a major difference between the way the *JSettlers* agents and human players negotiate is that people often make partial offers like example (1a). People typically do not make a complete trade offer immediately, but incrementally negotiate towards a more specific, and mutually acceptable, offer.

For this reason, we tested a **partialising agent** that initially computes a complete trade offer but then partialises it to only specify the resources it wants—i.e., it starts a negotiation with (1a), as opposed to (1b). It then reverts back to the original negotiation strategy. In this way, we test how our agents fare when adopting a human-like initial offering strategy.

Results and Discussion. Making an initial partial offer does not affect the agent's chances of winning and results in a small but significant increase in the number of trades it makes. But this agent does make many more trade offers than the original agents of which fewer are successful. The latter is due to the fact that agents cannot accept partial offers: typically, the complete offer that complements the initial partial offer was made by one of the original agents.

6.7 Memory: Forgetting Beliefs

The original agent and our modified agents so far do not have a realistic model of human memory: they do not forget any observed information. Therefore, we made the agent forget his beliefs about his opponents' resources after a certain time. For example, if the agent did not receive any new evidence about player-2's clay resources, the belief was set to 0 for the **forgetting pessimist proposer** and to 2 for the **forgetting cautious proposer** (so the cautious proposer assumes a favourable value in the absence of evidence to the contrary). In future work we will replace this by a more realistic and cognitively adequate memory model.

Because the *JSettlers* system does not maintain a global time, we used the number of messages sent by the *JSettlers* game server to approximate passing time. In a typical game, the server sends between 3500 and 5500 messages. We varied the time after which an agent forgets a belief from 5 to 1000. Note that ignorant agents are agents with a forgetting latency of 0.

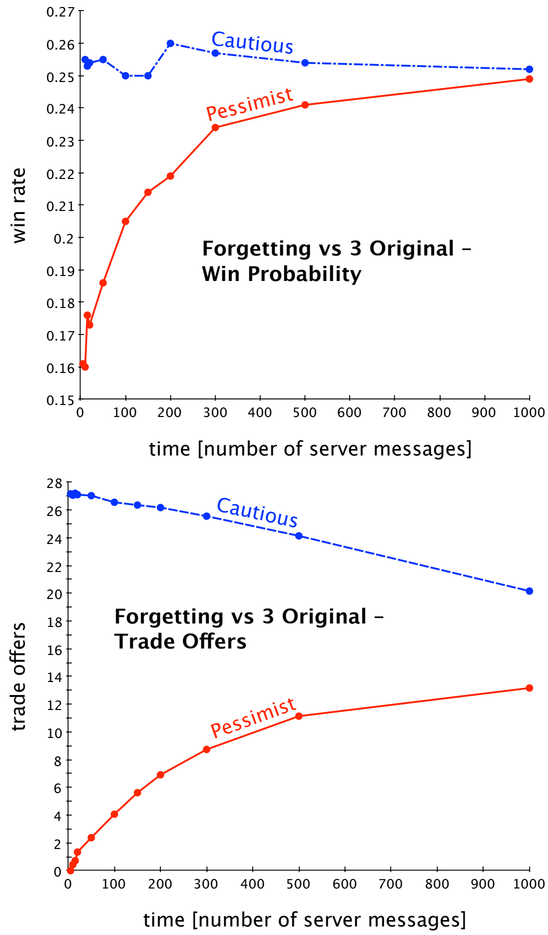

Figure 2: Win rates and number of offers over the latency before information is forgotten.

Results and Discussion. Figure 2 shows how the win rate and the number of trade offers changes with the time an agent can remember his beliefs. The cautious forgetting agent can, once more, compensate his lack of knowledge by assuming a desired outcome will occur in the absence of evidence to the contrary. And while the number of successful offers and total trades is constant for the cautious proposer and is the same as that of the original agent, the more forgetful he is, the more he makes unsuccessful trades offers and his negotiations become less effective.

The pessimistic forgetting agent makes fewer trade offers and wins fewer games the more forgetful he is. Only with a considerable ability to remember information (i.e. only if he forgets information after about 500 to 1000 time steps, which is already a considerable part of the entire game) does his performance approach that of the original agent.

Thus, only if a forgetful agent adopts favourable

	offers	success-ful offers	total trades	TIO	trades match-ing TIO	trades not match-ing TIO	successful offers/ offers
original	13.5	3.8	7.6	6.5	2.5	1.1	0.28
ignorant · cautious	30.6	3.8	7.6	10.2	1.8	1.9	0.12
omniscient · BP/resource	13.7	3.8	7.7	6.5	2.6	1.1	0.28
partialising initially	22.8	3.8	7.5	6.6	2.6	3.6	0.17
forgetting pessimist 200	6.7	2.8	5.6	4.0	2.1	0.6	0.42
Novice Corpus	*21.5*	*2.6*	*5.2*	*10.4*	*1.9*	*1.4*	*0.12*

Table 2: Results of 4 agents of the same type, and the corpus of novice players. TIO are turn-initial offers.

beliefs in the absence of actual information, can he compensate for his deficient memory.

6.8 A Comparison with Human Data

We now compare the simulated behaviour against the negotiating and trading behaviour of people playing *Settlers*. Our human data, shown in the last row of Table 2, is taken from an annotated corpus of humans playing *Settlers*, where detailed information about bargaining moves (offers, counteroffers, acceptance, rejection, etc.) and associated information about giveable and receivable resources that offers express are recorded (Afantenos et al., 2012).

As a first step, we used the annotated data from 7 games, where all players were new to the game, i.e. had not played *Settlers* before participating in the study. In future work, we will also compare the performance against players with varying degrees of expertise. As all players in the current set of games are members of a homogeneous population (novice *Settlers* players) we ran simulations where 4 agents *of the same type* play against each other. We only did this for the agents that are most interesting for a comparison (see Table 2).

Results and Discussion. The results from the simulations and the corpus given in Table 2 show that there is not one single agent from those we have investigated so far that exhibits the same behaviour as the novice players on all factors. The number of human trade offers is closest to the agent that initially partialises his trade offers; the successful trade offers and total trades are similar to the forgetting pessimist (with a forgetting latency of 200); and looking at the turn-initial offers (i.e. considering only the first offer of any given turn) and the ratio of successful offers over offers, the novices resemble the ignorant cautious (or the very similar ignorant optimistic) agent.

Thus, although novice players make very many trade offers, they are at the lower end of the spectrum when it comes to making successful offers, which in turn limits the number of trades they make. Note that the high number of trade offers is not simply a result of adopting a negotiation strategy of making partial offers (which, by definition, require at least one other offer to complete it): the high number of turn-initial offers shows that the number of game turns in which people try to trade is at the upper end of the spectrum.

These results are consistent with our previous explanations. People don't have a perfect memory (and for novice players, who have to keep track of many unfamiliar aspects of the game, this seems even more true). So they forget information. On the other hand they make many more attempts to trade, which is consistent with our suggestion that one strategy to compensate for an imperfect memory (up to a point) is to negotiate more.

7 Conclusions and Future Work

Strategic conversation, where the agents' goals diverge, is typically modelled as a game with a known analytic solution, where standard algorithms for identifying optimal actions apply. But human behaviour often diverges from game theoretic solutions and furthermore such models do not apply to noncooperative dialogues in the context of a complex game like *Settlers* that lacks any analytic solution. We, therefore, presented an alternative approach—an empirical framework in which on can evaluate how distinct dialogue strategies fare in the fact of distinct belief and memory models, including models that exhibit human-like errors like forgetting.

We first established that the existing *JSettlers* trading strategy correlates with winning. We then experimented with various models of the means for achieving such trades, and compared these models with the behaviour exhibited in a corpus of people playing *Settlers*.

Our agents varied on the extent to which they have accurate *beliefs* about hidden aspects of the game state, and the extent to which they can *remember evidence* for those hidden values. Agents with limited accuracy in their beliefs have three coping strategies for handling missing evidence: being *optimistic* and always assuming that the desired offer will have the desired outcome; being *cautious* (e.g., the original *JSettlers* agent) by using the available evidence to estimate hidden values but defaulting to favourable hidden values when evidence is inconclusive; and being *pessimistic* by assuming unfavourable values when evidence is inconclusive.

Belief and memory models affect the number and the quality of trade offers that agents make. An agent with perfect knowledge of the other agents' resources and a non-optimistic negotiation strategy makes offers that are much more likely to result in trades, but he fails to outperform his opponents' win rate, even though the opponents' belief models are more fallible. We speculate that this is because the opponents are more likely to enter into a negotiation (given that they default to assuming it will have the desired effect), and so increases the likelihood that they make an advantageous trade, if not the initially desired trade. Thus some ignorance can be bliss!

Indeed, being completely ignorant can be compensated for by being optimistic. This results in a much less efficient negotiator but more trades overall. On the other hand, capping the number of offers inhibits the advantages of this strategy. We can draw similar lessons from the forgetting agents, which show that (all else being equal) being able to remember more about your opponents' resources increases your chances to win. Or just be more optimistic (and less efficient).

We also took our first steps towards comparing our computational negotiation models with the strategies deployed by people playing *Settlers*. We showed that the predominant way of making partial offers observed in the human corpus does not change the win rate of our existing agents, and the observed changes in the number of offers and successful trade offers are consistent with our account. We also showed that while none of our agents directly models human performance, their performance is comparable in many ways. Of our agents, the ignorant optimistic agent is closest to the novice human performance exhibited in the corpus. In future work we will create cognitively more plausible agents and evaluate them by letting them play against humans.

While it was necessary for all agents to use the same build strategy (except for the random agent) so as to reduce the number of variables in our simulations, this is clearly an oversimplification in that the agents enjoy almost perfect predictions of their opponents' build plans.

In future work, we will explore the relative merits of revealing vs. concealing information about intentions when the agents all deploy distinct build strategies. We will also enhance the negotiation strategies by allowing agents to initiate a negotiation for a 'second best' trade when they believe their best trade won't have the desired effects, and we will investigate the benefits and costs of deception in a trade negotiation.

Acknowledgments: This work is supported by ERC grant 269427 (STAC). Thanks to Robin Hill for discussions on statistics and to Nicholas Asher, Oliver Lemon and 3 anonymous reviewers for comments and feedback.

References

S. Afantenos, N. Asher, F. Benamara, A. Cadilhac, C. Dégremont, P. Denis, M. Guhe, S. Keizer, A. Lascarides, O. Lemon, P. Muller, S. Paul, V. Popescu, V. Rieser, and L. Vieu. Modelling strategic conversation: model, annotation design and corpus. In *Proceedings of the 16th Workshop on the Semantics and Pragmatics of Dialogue (Seinedial)*, Paris, 2012.

D. Ariely. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. Harper-Collins, 2008.

Nicholas Asher and Alex Lascarides. Strategic conversation. *Sematics & Pragmatics*, 6(2):1–62, 2013.

D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

M. Pfeiffer. Machine learning applications in computer games. Master's thesis, Technical University of Graz, 2003.

V. Rieser, O. Lemon, and S. Keizer. Opponent modelling for optimising strategic dialogue. In *16th Workshop on the Semantics and Pragmatics of Dialogue (SeineDial'12)*, 2012.

G. Roelofs. Monte carlo tree search in a modern board game framework. research paper available at `umimaas.nl`, 2012.

L Savage. *The Foundations of Statistics*. John Wiley, 1954.

Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*. Cambridge University Press, 2009.

N. Sweeney. *Applying Machine Learning Techniques to an Imperfect Information Game*. PhD thesis, School of Computing, Dublin City University, 2012.

I. Szita, G. Chaslot, and P. Spronck. Monte-carlo tree search in Settlers of Catan. In H. van den Herik and P. Spronck, editors, *Advances in Computer Games*, pages 21–32. Springer, 2010.

K. Teuber. *Die Siedler von Catan: Regelheft.* Kosmos Verlag, Stuttgart, Germany, 1995.

R. Thomas. *Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic*. PhD thesis, Department of Computer Science, Northwestern University, 2003.

Adam Vogel, Christopher Potts, and Dan Jurafsky. Implicatures and nested beliefs in approximate Decentralized-POMDPs. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, August 2013. Association for Computational Linguistics.

Z. Yong and W. Xinlin. Intrinsic preferences, revealed preferences and bounded rational decisions. *Systems Research and Behavioural Science*, 2012. doi: 10.1002/sres.2120.