

Persuasion in Complex Games

Markus Guhe, Alex Lascarides

University of Edinburgh, School of Informatics

Informatics Forum, 10 Crichton St

Edinburgh EH8 9AB, Scotland

m.guhe@ed.ac.uk, alex@inf.ed.a.uk

Abstract

We study the power of persuasion in a game where each player’s own preferences over the negotiation’s outcomes are dynamic and uncertain. Our empirical set up supports evaluating individual aspects of the persuasion and reaction strategies in controlled ways. We show how this general method gives rise to domain-specific conclusions, in our case for *The Settlers of Catan*: e.g., the less scope there is for persuading during the game, the more one must ensure one gains an immediate benefit from it beyond the desired trade.

1 Introduction

In this paper we study persuasion in a non-cooperative setting, which Gricean (1975) maxims don’t account for (Asher and Lascarides, 2013). Within game theory, standard negotiation models ascribe each player complete and static knowledge of his own (intrinsic) preferences over the negotiation’s outcomes (e.g., Binmore (1998)). So the associated models of *persuasion* focus only on the persuader manipulating his opponents’ *beliefs* about which outcomes are likely (e.g., Rubinstein (2007)). For instance, during trading, the receiver of an offer to exchange wheat for clay might declare he has no wheat, and indeed be lying, so as to persuade his opponent to accept his counteroffer of ore for clay.

But if trading is a fraction of the action sequence in a complex game, then a player’s estimates of which next trade would enhance (or hinder) his chances to eventually win may be wrong. Persuasion then has higher stakes: there’s a new potential *payoff* in manipulating an opponent’s *preferences* over the next trade, not just his beliefs; but there’s also a new *risk* because the persuader’s deficient perception of the potential benefits of a particular trade may mean persuading backfires on him.

In addition, the persuader risks revealing information about his own intentions or preferences via the persuasion move.

Studies on manipulating an opponent’s trading preferences exist in argumentation theory (e.g., Amgoud and Vesic (2014)), but these models focus entirely on the logical structure of successful persuasion moves—i.e., moves where the recipient is persuaded and so changes his behaviour in the intended way. They don’t consider the persuading agent having a false perception of his own payoffs, and so don’t model the above risk of successfully persuading in a complex game.

Persuasion in complex games is commonplace. While interactions between businesses are often modelled via Markov Decision Processes, in reality the game tree isn’t surveyable because a player may make an offer that his opponent didn’t foresee as a possible move. Similarly, in board games like *Civilisation* and *The Settlers of Catan* (or *Settlers*) there are unbounded options for trading due to, for instance, the capacity to promise a specific future trade: e.g., *I’ll give clay for wheat now and ore when I get it if you don’t block me*. So standard algorithms for computing preferences over the outcomes of the current negotiation, like backwards induction and its variants (Shoham and Leyton-Brown, 2009), break down (Cadilhac et al., 2013).

We therefore need a general method for exploring the benefits and risks of persuasion in contexts that go beyond the ones modelled in standard negotiation games or argumentation theory. We supply a method here, using game simulations among computer agents whose symbolic strategies differ in transparent and controlled ways. We identify the following: when a persuasion move is likely to be successful (i.e., the recipient is persuaded); when successful persuasion results in a higher chance to win the overall game; and conversely when attempts to persuade are ineffective in improving win rates, even if they’re successful.

Our empirical set up provides a proof by demonstration that one can rapidly design, test and adapt symbolic persuasion strategies, with adaptation being guided by the quantitative performance metrics from game simulations. Specifically, we use our method to modify an existing agent that plays *Settlers*, and the result is a more effective player. Previous work on automatically learning *Settlers* strategies has shown that a decent prior player is critical for learning to succeed (Szita et al., 2010; Pfeiffer, 2003). We provide a principled way to build such priors, but investigating whether they enhance machine learning is future work.

In section 2 we describe related research, including work on agents that play *Settlers*. In section 3 we describe the rules of *Settlers* and the implemented agents that we use as a starting point. We then present our experiments, in which we manipulate the context in which the persuading agent chooses to perform a persuading move, the type of persuading move he attempts, and the strategies opponents adopt for accept or rejecting persuading moves. We provide quantitative metrics via game simulations of the effects of their different policies—e.g., their win rates and the number of persuasion moves executed. Our experiments radically discriminate among the persuasion strategies, identifying the strong strategies from the weak ones even though our game lacks any analytic solution.

2 Related Work

Negotiation in game theory (e.g., Binmore (1998), Brams (2003)) models when and how one suffers from the ‘winner’s curse’ (i.e., overpaying for an item, given the opponents’ preferences) and problems analogous to the prisoner’s dilemma (i.e., can one player trust the other to voluntarily cooperate during negotiation). But since each player has a complete and static model of his own preferences over the outcomes of negotiation the scope of persuasion gets restricted to persuading an opponent to change his beliefs but not his preferences over trades. Consequently game-theoretic models of persuasion (e.g., Rubinstein and Glazer (2006)) focus on the problem of predicting the *credibility* of the persuasive move. We address different questions: if one isn’t certain about which trades will help you, or hurt you, for winning the overall (complex) game, then how can one balance the benefits and risks of successfully persuading? And

hence at what stage in a complex game is successful persuasion most likely to increase one’s chances of winning the overall game? We propose an empirical method for answering these questions.

Our domain of study is *Settlers* (see section 3 for motivation). Empirical approaches to modelling *Settlers* deploy Monte Carlo Tree Search (Szita et al., 2010; Roelofs, 2012) and reinforcement learning (Pfeiffer, 2003). But even though they all use a simplified game, with no trading or negotiating, they all need a decent prior model for learning to succeed. So their priors encode sophisticated strategies, defined via complex *hand-coded heuristics*. Our work contributes to the general problem of developing decent priors: we supply an empirical framework where hand-coded heuristics can be rapidly designed and improved in light of quantitative performance metrics; e.g., Guhe and Lascarides (2013; 2014b) where we (a) identify negotiation strategies in *Settlers* that compensate for deficiencies in belief, e.g., memory loss, and (b) improve the building strategy used by our agents. Here, we identify effective persuasion strategies.

In trade negotiations, the persuading agent aims for either:

1. **More Trades:** i.e., a desired trade he might not achieve otherwise (e.g., *If you accept this trade, you’ll get clay and be able to build a road*); or
2. **Fewer Opponent Trades:** i.e., he stops two opponents from trading with each other (e.g., *Don’t trade with him! He’s about to win!*)

Kraus and Lehmann (1995) propose hand-built symbolic strategies for performing both these kinds of persuasion moves within the complex game *Diplomacy*, but the individual aspects of the strategies aren’t evaluated in controlled and transparent ways. We supply an empirical framework for doing just that. Here, we focus on game simulations for testing only those persuasion strategies that aim for More Trades; we address persuasion strategies aiming for Fewer Opponent Trades (FOT) in Guhe and Lascarides (2014a).

Achieving a successful persuasion move—i.e., one where the opponent is persuaded—is dependent on the persuading agent’s ability to adapt his persuasive argument to the current context and his

type of opponent. In a game of imperfect information, some executed persuasion moves are unsuccessful; i.e., they fail to persuade. So in this paper we explore how the persuading agent’s ability—or inability—to articulate arguments to opponents of various types should impact on his decisions about when to execute a persuading move so as to maximise his chances of winning the overall game.

3 The Settlers of Catan and JSettlers

The domain for our experiments is the board game *The Settlers of Catan* (or *Settlers*, (Teuber, 1995); www.catan.com). We chose it for its complexity: it is multi-player, partially observable, non-deterministic and dynamic; and further, with negotiations being conducted in natural language, the game’s options are unbounded (see earlier discussion). Thus our experiments prove that one can rapidly design and improve persuasion strategies in a principled and empirically grounded way even when game-theoretic algorithms for optimisation break down.

Settlers is a win–lose board game for 2 to 4 players. Each player acquires resources (ore, wood, wheat, clay, sheep) and uses them to build roads, settlements and cities on a board shown in Figure 1. This earns Victory Points (VPs); the first player with 10 VPs wins. Players can acquire resources via the dice roll that starts each turn and through trading with other players—so they negotiate trades. Players can also lose resources: e.g., a player who rolls a 7 can rob from another player. What’s robbed is hidden, so players are uncertain about their opponents’ resources. Deciding what resources to trade depends on what you want to build; e.g., a road requires 1 clay and 1 wood. Because *Settlers* is a game of imperfect information, agents frequently engage in ‘futile’ negotiations that result in no trade; i.e., they miscalculate the equilibria (Afantenos et al., 2012).

Our experiments modify an existing *Settlers* playing environment and automated *Settlers* player called *JSettlers* (`jsettlers2`, Thomas (2003)). *JSettlers* is a client–server system: a server maintains the game state and passes messages between the players’ clients, which can run on different computers. Clients can be humans or computer agents. Here, we report on simulations between computer agents.

The *JSettlers* agent goes through multiple phases after the dice roll that starts his turn:



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

1. Deal with game events: e.g. placing the robber; acquiring or discarding resources.
2. Determine legal and potential places to build.
3. Find the *Best Build Plan* (BBP), viz. the agent’s estimate of which build action gets him to 10 VPs in the shortest estimated time.
4. Try to execute the BBP, including negotiating and trading with other players.

Since we wish to study persuasion, our agents vary only in their policies for step 4, cf. section 4. Thomas (2003) describes steps 1–3. Here it only matters that the existing decisions on when to trade mean trading correlates with winning (Guhe and Lascarides, 2013).

In step 4 all agents have three existing possible responses to a trade offer: accept, reject or counteroffer. We equip our persuading agent with one more: to persuade an opponent to accept his trade offer. In our experiments, we vary the strategy for choosing among this expanded set of actions, and the strategies for reacting to the new option.

4 Evaluating Persuasion Moves

4.1 Motivation

There are a whole host of persuasive arguments that can accompany a trade offer—*Settlers* doesn’t restrict the types of trades nor the reasons for trading in any significant way. A small selection of possible persuasion moves is:

- (1) Give me 1 ore for 1 wheat and you can immediately build a settlement, which you can’t build without the wheat.

- (2) Give me 1 ore for 1 wheat and only then will you have enough wheat to make a trade with your 3:1 port.
- (3) If you give me 1 ore for 1 wheat, you can use the wheat to trade for James' clay, so that you can build your road.
- (4) If you give me 1 ore for 1 wheat, I won't rob you the next time I'm playing a knight card.
- (5) If you give me 1 ore for 1 wheat, I'll build a road that blocks Nick from that port.

So the benefits and risks of persuasion will depend on (at least):

\mathcal{P} 's ingenuity: the range of contexts where the persuading agent (who we'll label \mathcal{P}) can articulate a persuasive move like those in (1) to (5) and beyond.

\mathcal{P} 's caution: In those contexts where his ingenuity provides a candidate persuasion move, \mathcal{P} 's strategy for deciding whether to actually make that move; and

\mathcal{G} 's gullibility: how inclined the recipient (labelled \mathcal{G}) is to accept \mathcal{P} 's persuasion move and hence also the trade offer.

Ingenuity and caution are distinct factors that determine a persuader's player type: ingenuity affects the persuader's range of options (he is more or less able to generate a candidate persuasion move); caution affects the persuader's penchant for actually executing a persuasion move when such a move is an option. Our experiments vary both factors, because the optimal level of caution may be different for an ingenious vs. non-ingenious agent—after all, an ingenious cautious player's behaviour is not in general equivalent to that of a non-ingenious, non-cautious player.

Asher and Lascarides (2013) show that a rational \mathcal{G} will normally accept \mathcal{P} 's speech act—and a persuading move in particular—if \mathcal{G} believes \mathcal{P} to be *sincere* (i.e., \mathcal{P} believes what he says) and *competent* (i.e., what \mathcal{P} believes is true). But \mathcal{P} can appear sincere and competent without actually being so. For instance, \mathcal{P} can utter (1) but be ignorant about whether \mathcal{G} has the other resources he needs for a settlement (i.e., clay, wood and sheep) and/or he may lack evidence that building a settlement is better for \mathcal{G} than \mathcal{G} 's current build plan (whatever that is). In this case, \mathcal{P} is neither sincere nor competent. But even if \mathcal{G} lacks clay, wood and

sheep, it's still consistent for him to assume that \mathcal{P} was sincere (but inconsistent to assume he's competent), for \mathcal{G} 's resources aren't observable to \mathcal{P} and \mathcal{P} 's beliefs aren't observable to \mathcal{G} . Further, if \mathcal{G} does have clay, wood and sheep, then because \mathcal{G} is uncertain about his own relative preferences over build plans, it's consistent for \mathcal{G} to assume that \mathcal{P} is both sincere and competent in (1)'s implicated content, that building a settlement is both possible *and* better for \mathcal{G} . Thus there's scope for \mathcal{P} to successfully bluff, getting \mathcal{G} to accept his persuasion move even though he's neither sincere nor competent. Our experiments thus investigate when bluffing succeeds, and whether successfully bluffing helps \mathcal{P} win the overall game.

4.2 The Agents' Contexts

We start with a persuading agent \mathcal{P} with *maximal ingenuity*—i.e., he can make a persuasion move every time he makes a trade offer and is unrestricted in the number of such moves he can make in the course of the game. Further, we make \mathcal{G} *maximally gullible*: he assumes \mathcal{P} 's persuasion move is convincing so long as the proposed trade is executable. We then vary \mathcal{P} 's *caution*, by making \mathcal{P} start executing persuasion moves only once the first agent reaches a specified number of VPs. We call this factor *VP*. In Guhe and Lascarides (2014a) we showed that the timing of persuasion moves is crucial and moves early and late in the game are much less effective than if they are used when the first player has reached around 7 VPs.

We call these agents *simple*. In terms of Guhe and Lascarides (2013) these agents are both *ignorant*, in that they use only observable information (VPs for \mathcal{P} , his own resources for \mathcal{G}) to decide what to do. A simple \mathcal{P} is also relatively incautious, because the leader's VPs is the only factor that prevents \mathcal{P} 's trade offer from having an accompanying persuasion move too.

From this starting point, we will then vary \mathcal{P} 's degree of *caution*, by restricting the contexts (over and above *VP*) in which \mathcal{P} actually chooses to make a persuasion move, and \mathcal{G} 's *gullibility* by restricting the contexts in which \mathcal{G} accepts \mathcal{P} 's persuasion moves.

4.3 Method for Simulation and Analysis

A simulation for testing the different persuasion moves consists of 1 persuading agent (\mathcal{P}) playing 3 non-persuading opponents (\mathcal{G}) 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 run on the same machine, and 10,000 games take about 1 hour on a current desktop computer.

Apart from the agents' win rates, we measure how many persuasion moves \mathcal{P} actually makes: the fewer persuasion moves \mathcal{P} needs to gain a significant advantage in winning, the more efficient they are in achieving desirable effects.

We performed Z-tests with $p < 0.01$ to test significance of win rates against the null hypothesis. This means that win rates between 0.24 and 0.26 don't differ significantly from the null-hypothesis; so we highlight the 0.26 threshold in the graphs below. We report the average numbers for \mathcal{P} for each simulation, and averages across all three of \mathcal{P} 's opponents. Due to the large number of games per simulation even small differences can be significant. At the same time, there were no significant differences between the three \mathcal{G} s. Persuasion does not affect the average length of the game, which is consistently between 21 and 21.5 rounds.

5 Simple vs. Cautious \mathcal{P}

5.1 $\mathcal{P}:\emptyset$ vs. $\mathcal{P}:PB$

In the first set of simulations we compared simple agents (i.e. agents that make/accept the maximum number of persuasion moves) and then restricted \mathcal{P} to a more self-serving context:

1. **None (\emptyset):** \mathcal{P} using this context makes a persuasion move with every trade offer *proviso* the *VP* factor; \mathcal{G} using this context accepts all persuasion moves and the accompanying trade offer if the trade is executable (i.e. \mathcal{G} has the resources for making the trade).
2. **Persuader Build (PB):** \mathcal{P} makes the persuasion move iff *VP* is satisfied *and* the proposed trade allows him to build immediately, i.e. to execute his BBP after making the trade.

A \mathcal{P} who adopts *PB* is relatively cautious: he's attempting to mitigate the risk of his deficient preferences over trades by ensuring that all successful persuasions result not only in his desired trade but also in the immediate benefit of building.

Figure 2 shows the simulation results for the configuration ($\mathcal{P}:\emptyset, \mathcal{G}:\emptyset$)—i.e. \mathcal{P} can make an unlimited number of persuasion moves ($N=\infty$) and \mathcal{G} accepts all such moves—as well as for the configuration ($\mathcal{P}:PB, \mathcal{G}:\emptyset$). $\mathcal{P}:PB$'s win rates are al-

most as good as $\mathcal{P}:\emptyset$'s (0.363 vs. 0.377 at 2 VPs; 0.274 vs. 0.285 at 8 VPs) but he needs substantially fewer persuasion moves for this (15.4 vs. 40.8 at 2 VPs; 1.4 vs. 6.0 at 8 VPs).

Realistically, a fully ingenious \mathcal{P} risks irritating his opponents and making them suspicious if he makes a persuasion move every time he can—even 15 moves in the course of a game (cf. $\mathcal{P}:\{PB, VP=2\}$) is more than humans do according to our corpus data (Afantenos et al., 2012). Figure 3 shows what happens if the number of persuasion moves \mathcal{P} can make are limited ($\mathcal{P}:\{N \in [1, 3]\}, \mathcal{G}:\emptyset$). $\mathcal{P}:PB$ achieves a significant improvement over the null-hypothesis even when he only makes 1 move at most, so long as he makes that move after the first player reaches 6 VPs. (This is consistent with our results in Guhe and Lascarides (2014a).) The less cautious $\mathcal{P}:\emptyset$ needs to be able to make at least 3 moves to gain a significant advantage. The right graph in Figure 3—depicting the number of moves \mathcal{P} actually made—also shows that even though $\mathcal{P}:PB$ makes fewer moves than $\mathcal{P}:\emptyset$, he achieves a much higher win rate. So perhaps surprisingly, the less ingenious \mathcal{P} needs to be more cautious.

In the following, we will only report on simulations where \mathcal{P} can make an unlimited number of persuasion moves (i.e., $N = \infty$), because the main effect for N is the same across simulations: the higher N is, the more moves \mathcal{P} makes and the more games he wins.

5.2 Number of gullible agents

An agent's success is always highly dependent on his opponents. So we checked how much \mathcal{P} 's performance depends on the number of persuadable opponents he plays against. These simulations vary the number of \mathcal{G} opponents who accept persuasion moves vs. those (non- \mathcal{G}) opponents who never accept them. For conditions ($\mathcal{P}:\emptyset, \mathcal{G}:\emptyset$) and for ($\mathcal{P}:PB, \mathcal{G}:\emptyset$), \mathcal{P} retained a big advantage over all three opponents even when only one of them is persuadable. Further, deploying *PB* helps \mathcal{P} achieve almost the same win rate as without it, but with fewer than half of the persuasion moves.

| config. | wins | | moves made | |
|-----------------|-------------|-----------|-------------|-----------|
| | \emptyset | <i>PB</i> | \emptyset | <i>PB</i> |
| 3 \mathcal{G} | 0.383 | 0.363 | 40.7 | 15.4 |
| 2 \mathcal{G} | 0.342 | 0.341 | 47.6 | 19.1 |
| 1 \mathcal{G} | 0.302 | 0.315 | 60.3 | 24.7 |

The reason why the persuader needs more moves the fewer opponents are gullible is that




Figure 2: Win rate and persuasion actually moves made, against the VP factor (i.e., the leader’s minimum VPs before persuading can start). The dashed line is $\mathcal{P}:\emptyset$, the solid line is $\mathcal{P}:PB$.




Figure 3: Win rate and moves made against the VP factor. Dashed lines are $\mathcal{P}:\emptyset$ and solid lines $\mathcal{P}:PB$.

more of the persuader’s trade offers are unsuccessful (non-gullible agents accept offers at a normal rate). So \mathcal{P} has to make more offers to get the trades he wants.

6 A more discerning \mathcal{G}

So far, our \mathcal{G} agents are so gullible that they don’t test the persuasive argument for sincerity or competence. We now restrict \mathcal{G} ’s gullibility: instead of accepting all persuasion moves where the trade is executable ($\mathcal{G}:\emptyset$), we make \mathcal{G} accept whatever \mathcal{P} ’s persuasive move is only if \mathcal{G} can build something or make a bank/port trade as a result of trading (in the following we abbreviate *trade with the bank or an available port to bank trade*). In other words, factors for \mathcal{G} accepting a persuasion move are:

1. **Gullible Build (GB):** \mathcal{G} accepts the persuasion move only if it enables him to build a type of piece that he cannot build without making the trade.
2. **Gullible Bank Trade (GBT):** \mathcal{G} accepts the

persuasion move only if after making the trade he can make a bank trade immediately.

3. **GBoBT:** The disjunction of these two cases.

Note that $\mathcal{G}:GB$ by default assumes that \mathcal{P} is sincere and competent on persuasive moves like (1), and $\mathcal{G}:GBT$ by default assumes that \mathcal{P} is sincere and competent on persuasive moves like (2).

Here it is important to distinguish the *persuasion move* from the *trade offer* that it is accompanying: Even if \mathcal{G} does not accept the persuasion argument (e.g., \mathcal{G} infers \mathcal{P} ’s persuasion argument is not competent), he will still evaluate the trade offer in it’s own right. For example, in (1), \mathcal{G} may still agree to exchange 1 ore for 1 wheat, even if this does not enable him to immediately build the settlement as \mathcal{P} claims. That is, \mathcal{G} never rejects a trade offer with a persuasion move if he would have accepted it without the persuasion.

Figure 4 gives the results for both $\mathcal{P}:\emptyset$ and $\mathcal{P}:PB$. In all cases, \mathcal{P} fares better in the PB context than in the \emptyset one and with fewer persuasion moves; i.e., $\mathcal{P}:PB$ is not only more effective but

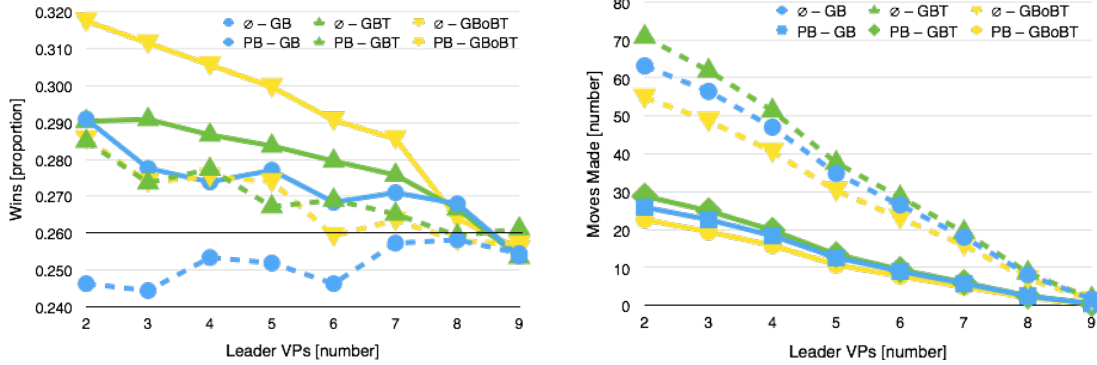


Figure 4: Win rate and moves made when varying \mathcal{G} 's gullibility. Dashed lines are $\mathcal{P}:\emptyset$; solid lines $\mathcal{P}:PB$.

also more efficient. Indeed, there is no effect for $(\mathcal{P}:\emptyset, \mathcal{G}:GB)$ —i.e., no agent gains an advantage—but there is an effect for $(\mathcal{P}:PB, \mathcal{G}:GB)$. So as \mathcal{G} gets less gullible, \mathcal{P} should get more cautious (i.e., play with strategy PB). On the other hand, as we saw in sections 5.1 and 5.2, so long as at least one of \mathcal{P} 's opponents is *maximally gullible* ($\mathcal{G}:\emptyset$), \mathcal{P} should be maximally incautious (i.e., $\mathcal{P}:\emptyset$).

For $\mathcal{G}:GBT$, \mathcal{P} gains an advantage for both of his contexts. Thus, the potential benefit for \mathcal{G} if \mathcal{P} makes move (2) is smaller than that for move (1); conversely, \mathcal{P} 's risk in successfully making move (1) is smaller than that for move (2). When \mathcal{G} accepts persuasion attempts of both kinds so long as it's consistent with sincerity and competence—i.e. $\mathcal{G}:GBoBT$ —then \mathcal{P} has a bigger advantage than if \mathcal{G} uses only one of the contexts, and \mathcal{P} needs even fewer moves. So, one general observation here is that the more types of persuasion moves \mathcal{G} accepts, the more successful \mathcal{P} is and the fewer moves \mathcal{P} needs to achieve this.

7 \mathcal{P} taking \mathcal{G} 's Context into Account

So far, \mathcal{P} does not reason about \mathcal{G} 's likely reaction when deciding whether to make a persuasion move. But as we said earlier, persuasion must appear sincere and competent to a rational \mathcal{G} to be successful. And \mathcal{P} can reduce the risk of miscalculating equilibria and making futile moves by reasoning about \mathcal{G} 's likely reaction. We investigate this by restricting \mathcal{P} 's ingenuity—he can only articulate moves of the form (1) or (2)—and \mathcal{P} 's caution in the following ways:

1. **Persuader Opponent Build (POB):** \mathcal{P} only makes a persuasion move only if he believes that it allows \mathcal{G} to build something that he cannot build without making the trade.

2. **Persuader Opponent Bank Trade (POBT):** \mathcal{P} makes the persuasion move only if \mathcal{P} believes that after making the trade, \mathcal{G} can immediately make a bank trade that he cannot make without the trade.

3. **POBoBT:** The disjunction of these cases.

Whether \mathcal{P} executes a persuasion move now depends on \mathcal{P} 's beliefs about \mathcal{G} 's resources. For instance, agent $\mathcal{P}:POB$ must believe that the resources \mathcal{G} gets in the proposed trade are necessary and sufficient for \mathcal{G} to immediately build. So $\mathcal{P}:POB$ is in effect only making persuasion moves of form (1), and executes such a move only if \mathcal{P} believes that a \mathcal{G} player of the following type will accept it: (a) \mathcal{G} is rational, and so accepts a move iff \mathcal{G} believes it's sincere and competent; and (b) \mathcal{G} defaults to believing moves are sincere and competent. Similarly, $\mathcal{P}:POBT$ only makes persuasion moves of the form (2) and only executes them if \mathcal{P} believes a \mathcal{G} player of the above type will accept it; $\mathcal{P}:POBoBT$ is slightly more ingenious, using persuasion moves of both types.

The agents use the standard *JSettlers* belief model, i.e. no memory loss and fully accurate beliefs about how many resources each opponent has, but some are of unknown type because of robbing. In terms of Guhe and Lascarides (2013), \mathcal{P} is relatively cautious: he does not take \mathcal{G} 's unknown resources into account, i.e. he only makes a persuasion move, if he *knows* that \mathcal{G} can execute the build or bank trade he promises— \mathcal{P} does not bluff.

Depending on his gullibility configuration, \mathcal{G} accepts different persuasion arguments, e.g. $\mathcal{G}:GB$ is only susceptible to the arguments of $\mathcal{P}:POB$ (or, $\mathcal{P}:POBoBT$) but not $\mathcal{P}:POBT$.

Similar to the previous result, $\mathcal{P}:POB$ does not improve his win rate but $\mathcal{P}:\{PB, POB\}$ does. And




Figure 5: Win rate and moves made when \mathcal{P} takes \mathcal{G} 's context into account. Dashed lines are $\mathcal{P}:\emptyset$; solid lines $\mathcal{P}:PB$.

when being more selective about making persuasion moves (by adopting PB), adding POB does not reduce \mathcal{P} 's win rate, but he needs only about a quarter of the moves.

In the $(\mathcal{P}:\{PB, POBT\}, \mathcal{G}:GBT)$ context, \mathcal{P} 's win rate is similar to the one without \mathcal{P} taking \mathcal{G} 's context into account. This strategy is more efficient for \mathcal{P} (fewer moves) and more effective (higher win rate) than POB . Again, the PB context is more effective and efficient than the \emptyset context.

Finally, in $(\mathcal{P}:\{PB, POBoBT\}, \mathcal{G}:GBoBT)$ \mathcal{P} makes both kinds of persuasive moves as well as both kinds of assessments about \mathcal{G} 's state, and \mathcal{G} is selective about both types of moves. The added opportunities that \mathcal{P} obtains through his increased ingenuity compared to an agent who can make only one type of argument leads to more persuasive moves being executed and a higher win rate.

Comparing these results to the simulations when \mathcal{P} does not take \mathcal{G} 's gullibility into account, we again see that \mathcal{P} can increase its efficiency (he makes fewer moves) without sacrificing his effectiveness (the win rates do not differ substantially).

8 Conclusions

In this paper we used *The Settlers of Catan* to investigate the power of persuasion in a multi-player, partially observable, non-deterministic, dynamic, unbounded game. We established an empirical method involving game simulations, with the heuristics that the persuading agent and his recipients use being evaluated in controlled ways and improved upon.

We found that the more ingenuity the persuader has at articulating persuasive arguments, and the more gullible his recipients, the more successful he becomes at winning the overall game. Indeed,

one gullible agent is sufficient for the persuader to gain an advantage over all three opponents. If he lacks ingenuity and so is restricted to only certain kinds of arguments, then it helps to make performing a persuasive move dependent on whether the proposed trade will enable him to immediately build. The persuader can also increase the proportion of his persuasion moves that are successful without harming his win rate by reasoning about how his opponent will react.

Gullible agents, who assume the persuader is sincere and competent by default, cannot improve over the null-hypothesis—a 25% win-rate. But they ‘lose less’ if they are selective about the persuasion moves they comply with; here, if you comply with just one kind of persuasion move, it should be the one like (1) (i.e., you can immediately build but only if you execute the proposed trade).

We are currently collecting data on how persuasive human opponents find the More Trade persuasion moves we investigated here and will then investigate persuasion that aims for Fewer Opponent Trades. We will then use our *Settlers* environment to test our best persuasive agents against humans. We will also investigate the impact of our improved priors on automatically learning *Settlers* strategies and opponent modelling similar to the work by Gal et al. (2004) in order to adapt to human opponents over the course of a game.

Acknowledgments

We would like to thank the reviewers for helpful comments. This work is funded by ERC grant 269427 (STAC).

References

- Stergos Afantenos, Nicholas Asher, Farah Benamara, Anaïs Cadilhac, Cedric Dégremont, Pascal Denis, Markus Guhe, Simon Keizer, Alex Lascarides, Oliver Lemon, Philippe Muller, Soumya Paul, Vladimir Popescu, Verena Rieser, and Laure Vieu. 2012. Modelling strategic conversation: model, annotation design and corpus. In *Proceedings of the 16th Workshop on the Semantics and Pragmatics of Dialogue (Seinedial)*, Paris.
- Leila Amgoud and Srdjan Vesic. 2014. Rich preference-based argumentation frameworks. *International Journal of Approximate Reasoning*, 55.
- Nicholas Asher and Alex Lascarides. 2013. Strategic conversation. *Semantics & Pragmatics*, 6(2):1–62.
- Ken Binmore. 1998. *Game Theory and the Social Contract: Just Playing*. MIT Press.
- Steven Brams. 2003. *Negotiation Games: Applying game theory to bargaining and arbitration*. Routledge.
- Anaïs Cadilhac, Nicholas Asher, Alex Lascarides, and Farah Benamara. 2013. Preference change. submitted.
- Ya’akov Gal, Avi Pfeffer, Francesca Marzo, and Barabara J Grosz. 2004. Learning social preferences in games. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI), San Jose, California, July 2004*.
- H. Paul Grice. 1975. Logic and conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics Volume 3: Speech Acts*, pages 41–58. Academic Press.
- Markus Guhe and Alex Lascarides. 2013. Effects of belief and memory on strategic negotiation. In *Proceedings of SEMDIAL 2013 – The 17th Workshop on the Semantics and Pragmatics of Dialogue, Amsterdam, 16–18 December 2013*, pages 82–91.
- Markus Guhe and Alex Lascarides. 2014a. The effectiveness of persuasion in The Settlers of Catan. In *Proceedings of CIG 2014, The IEEE Conference on Computational Intelligence and Games*.
- Markus Guhe and Alex Lascarides. 2014b. Game strategies for The Settlers of Catan. In *Proceedings of CIG 2014, The IEEE Conference on Computational Intelligence and Games*.
- Sarit Kraus and Daniel Lehmann. 1995. Designing and building a negotiating automated agent. *Computational Intelligence*, 11(1):132–171.
- Michael Pfeiffer. 2003. Machine learning applications in computer games. Master’s thesis, Technical University of Graz.
- Gijs-Jan Roelofs. 2012. Monte Carlo Tree Search in a modern board game framework. Research paper available at umimaas.nl.
- Ariel Rubinstein and Jacob Glazer. 2006. A study in the pragmatics of persuasion: a game theoretical approach. *Theoretical Economics*, 1(4):395–410.
- Ariel Rubinstein. 2007. What is the proper role of game theory to other disciplines? In Vincent F Hendricks and Pelle Guldberg Hansen, editors, *Game Theory: 5 Questions*. Automatic Press.
- Yoav Shoham and Kevin Leyton-Brown. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*. Cambridge University Press.
- István Szita, Guillaume Chaslot, and Pieter Spronck. 2010. 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.
- Klaus Teuber. 1995. *Die Siedler von Catan: Regelheft*. Kosmos Verlag, Stuttgart, Germany.
- Robert S Thomas. 2003. *Real-time Decision Making for Adversarial Environments Using a Plan-based Heuristic*. Ph.D. thesis, Northwestern University.