



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The effectiveness of persuasion in The Settlers of Catan

Citation for published version:

Guhe, M & Lascarides, A 2014, The effectiveness of persuasion in The Settlers of Catan. in *Computational Intelligence and Games (CIG), 2014 IEEE Conference on*. Institute of Electrical and Electronics Engineers, pp. 1-8. <https://doi.org/10.1109/CIG.2014.6932861>

Digital Object Identifier (DOI):

[10.1109/CIG.2014.6932861](https://doi.org/10.1109/CIG.2014.6932861)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Computational Intelligence and Games (CIG), 2014 IEEE Conference on

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



The Effectiveness of Persuasion in *The Settlers of Catan*

Markus Guhe, Alex Lascarides

School of Informatics
University of Edinburgh
Edinburgh EH8 9AB
Scotland

Email: m.guhe@ed.ac.uk, alex@inf.ed.ac.uk

Abstract—We present a method for obtaining a useful symbolic model of persuasion in a complex game, where players’ preferences over the outcomes of negotiations over resources are incomplete, uncertain, and dynamic. We focus on the problem of identifying the stage in the game where successfully persuading an agent to perform a particular action will have the most impact on one’s chances to win. Our approach exploits empirical data from game simulations, where the set up allows us to investigate individual aspects of the policy in suitably controlled ways. We demonstrate the effectiveness of the approach within the domain of *The Settlers of Catan* and present some specific lessons that can be learned for this particular game, e.g. that a tipping point in the game occurs for persuasion moves that are made when the leading player reaches 7 Victory Points.

I. INTRODUCTION

Strategic negotiation is a type of non-cooperative conversation, which the Gricean view of cognitive agents doesn’t account for [1]. Within game theory, standard models of negotiation assume that each player’s (intrinsic) preferences over the negotiation’s outcomes are static, complete and predefined (e.g., [2]). A player can be uncertain about his opponent’s preferences over the outcomes, but the preferences themselves don’t change. Thus models of persuasion have until now focussed entirely on the speaker revealing information about the current state so as to manipulate his opponents’ *beliefs* about which negotiation outcomes are likely (e.g., [3]). For instance, in a trading negotiation, 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 rock for clay.

But what if persuasion happens in an extensive game where trading is a small fraction of the overall action sequence? Further, what if the game is so complex that no player can be certain about even his own preferences over the next trade? Such scenarios fail to meet the informational demands of the definition of a standard negotiation game: preferences over trades are no longer static, complete or predefined. Thus standard models of persuasion don’t apply on their own either. But intuitively, the stakes of persuasion are higher: on the one hand one can now persuade a player to revise his *preferences* over trades, and not just his beliefs about which trades are likely; on the other hand, a persuading agent risks inadvertently persuading opponents to behave in ways that hurt him. A further risk arises in a time critical situation. Calculating whether a persuasion move will have its desired

immediate effect—i.e., changing the receiver’s behaviour—is computationally (or cognitively) expensive [3]. But if other negotiating players can make improved trade offers at any time in the negotiation, then the time taken to identify a convincing persuasive move may backfire.

Persuasion in such complex scenarios is commonplace: while ongoing business transactions between large commercial organisations are often modelled as some form of Markov Decision Process, in reality the game tree isn’t surveyable: e.g., a player may make an offer that his opponent didn’t foresee as a possible move in the game. Similarly, in board games like *Diplomacy* or *The Settlers of Catan* (or *Settlers*, the game we investigate here), the game is complex and includes unbounded options for trading (thanks to, for instance, the capacity to promise a specific future trade; e.g., *If you trade clay for wheat now, I’ll give you rock when I get it*). So standard algorithms for computing (extrinsic) preferences over the outcomes of the current negotiation, such as backwards induction and its variants [4], break down [5].

What’s needed for modelling persuasion in these complex games is a general method for exploring its benefits and risks that doesn’t require the information demanded by standard negotiation games, such as complete and static preferences over the negotiation’s outcomes. We explore one such general approach here, grounded in empirical data. Specifically, we provide a proof by demonstration that one can rapidly design, test and adapt symbolic persuasion strategies, with adaptation being guided by game simulations, whereby we achieve a rapid improvement in the agent’s ability to win the game.

We focus here on a particular issue: at what stage in the game would performing a successful persuasion, such as we find in a corpus of people playing the game [8], yield the most radical improvement to one’s chances of winning? This is a critical question, given the above risks in performing complex calculations to identify a convincing persuasion move. In effect, we offer an empirical method to identify the upper bound on the benefits of persuasion in a complex game.

In trade negotiations, the persuading agent aims for either:

- 1) **More Trades:** i.e., a desired trade he might not achieve otherwise (e.g., *But 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!*)

Our objective is to identify empirically the stage in the game when a successful persuasion of both these types results in the biggest increase to one’s chances of winning the overall game. We use game simulations in which 4 agents play *Settlers*, with one player’s persuasion strategy being different from the other three. We vary the *context* in which the persuading agent decides to make a persuading move, and vary also whether his persuasion move aims for the outcome MT (More Trades for himself) or FOT (Fewer Opponent Trades). Of course, these outcomes are desirable only if the agent’s policies on trading ensure that winning correlates with how many trades he executes (and likewise inversely correlates with how many trades his opponents perform). In earlier work, we showed these correlations hold for our agents [6], [7]. Here, we aim to identify which negotiation within a *sequence* of negotiations (and other actions) is the most critical for gaining an overall advantage in the game.

Because we are investigating at what point in a game a successful persuasion move would reap the most reward, we ignore the complex issue of *whether* the persuasion move is going to have the desired immediate outcome: i.e., whether the intention to revise the opponents’ behaviours gets realised. So in our experiments, we make the 3 ‘baseline’ agents against which the persuading agent plays maximally gullible: they are always persuaded. In [9] we report on simulations where the baseline agents are more selective about when they are persuaded; likewise, the persuading agent is endowed with a more restricted range of persuasive arguments that he can deploy, and so they’re appropriate in only a restricted set of contexts. But in this paper, the only real choice (and risk) the persuading agent deliberates on is the *optimal timing* in the course of the game of a successful persuasion move.

After discussing related work in section II, we describe the rules of *Settlers* in section III and the implemented agents and their strategies in section IV. In the remainder of the paper we then present our experiments, in which we manipulate the timing and type of persuasion move that an agent performs, and we provide quantitative metrics of the effects of these different behaviours in terms of win rates, the number of trades agents achieve, the number of resources they receive via trading, and also the number of opportunities that an agent had to actually choose a persuasion move, given his policies for when to perform it. Given that our experiments radically discriminate among the performances of agents with different persuasion strategies, this work is a proof by demonstration that one can rapidly design effective heuristics of game play, in an empirical and principled fashion, for complex games where standard algorithms for computing optimal action all break down.

II. RELATED WORK

Game theory as it applies to negotiation is well studied (e.g., [2], [10], [11]), with researchers observing when and how one can suffer from the ‘winner’s curse’ (i.e., overpaying for an item, given the opponents preferences) and when and how situations analogous to the prisoner’s dilemma occur (i.e., can one player trust the other to voluntarily cooperate during negotiation). In almost all these studies, the negotiation is a ‘one off’ game: they assume that each player has a complete and static model of his own preferences over the end states of

the negotiation. This restricts the scope of what persuasion can achieve: one cannot persuade an opponent that he has the wrong intrinsic preferences, only that his beliefs about the current state and/or the outcome of actions need revising.

Existing game-theoretic models of persuasion (e.g., [12]) assume that the persuading agent can access information about the current state that his opponents can’t, and moreover he has complete information about his own utility of a successful outcome of his persuasion move (i.e., the outcome where the opponent is persuaded to change his behaviour in the way the persuading agent intended). This is largely because this work focusses on predictive models of the *credibility* of the information declared in the persuasion. In contrast, we address a different problem: in a game where the benefit of a successful persuasion is uncertain, how can one weigh the risk of expounding effort to identify a convincing persuasion against its potential benefit? And in particular, at what stage in the game is a successful persuasion most likely to reap the most reward? We propose an empirical method for answering this question. Our study uses *Settlers* as a domain in which players cannot accurately calculate the utility of successfully persuading when they make their persuasive move.

There are several empirical approaches to modelling *Settlers*, but none of them includes trading or negotiating. Specifically, [13] and [14] use Monte Carlo Tree Search and [15] uses reinforcement learning. In addition to using simplified versions of the game problem, both approaches also demonstrate that a decent prior model is critical to the learning process being successful. They all use as a starting point players with already sophisticated strategies, defined via complex *hand-coded heuristics*. This paper contributes to the general problem of developing decent priors that make learning improved strategies over complex games a possibility, by supplying an empirical framework where hand-coded heuristics for playing the game can be rapidly designed and improved in light of their effects on game performance. In [6] we used this framework to identify which negotiation strategies in *Settlers* can compensate for deficiencies in belief, e.g., via memory loss, and in [7] we used it to improve the preference function over building options. Here, we use it to identify the point in the game where a successful persuasion move will have the most beneficial impact on trading and winning.

III. THE SETTLERS OF CATAN

Our example domain for investigating persuasion is the board game *The Settlers of Catan* (or *Settlers*, [16]; see www.catan.com). This game exhibits maximum complexity: it is multi-player, partially observable, non-deterministic and dynamic (thanks to dice rolls); and further, with negotiations being conducted in natural language, the game’s options are unbounded (see earlier discussion). This complexity makes *Settlers* an ideal domain for our purposes, with prior work showing that learning policies for playing *Settlers* are feasible *only if* the learning process starts with a very decent prior strategy [13], [14]. The work we report here is a proof by demonstration that one can rapidly design and evaluate such prior models in a principled and empirically grounded way.

Settlers is a win-lose board game for 2 to 4 players. Each player acquires resources (ore, wood, wheat, clay, sheep) and



Fig. 1. A game of *The Settlers of Catan* in JSettlers.

uses them in different combinations to build roads, settlements and cities on a board like the one 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 players converse to negotiate trades. 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 [17]. Players are free to agree any trade (any quantity or combination) with the exception of just giving resources away. Trades can also be linked to other game actions (*Make this trade with me, and I won’t rob you for 2 rounds*). 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 that result in no trade (i.e., they miscalculate the equilibria).

IV. PLANNING AND NEGOTIATION IN JSETTLERS

We build on an open source implementation called JSettlers (jsettlers2.sourceforge.net, [17]). 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 human players or computer agents, cf. also [6]–[8]. Here, we report on simulations between computer agents only.

The JSettlers agent goes through multiple phases after the dice roll that starts his 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 VPs 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.

As we are exploring an aspect of negotiation, all our agents adopt the same policy for doing steps 1–4, while different agents adopt different policies for doing step 5. [6] and [17] describe the policy for steps 1–4 in detail, but for the purposes of this paper the issues of when, how and why an agent decides that trading is its best available action are not important. For our purposes, what matters is that the existing implemented policy on trading is effective, in that on average, trading correlates with winning. In [6] we show this is the case for all our agents.

In this paper, we explore the effects of persuasion by varying step 5. All our agents have three existing possible responses to a trade offer: accept, reject or counteroffer. We equip our persuading agent with three more responses:

- 1) Issue a **trade embargo** against a player (as might result from *Don’t trade with him, he’s about to win*); an embargoed player will not be able to make any trades,
- 2) **Block** for a specified number of turns the offering agent from making both the specific trade that he’s just offered and any trade where he receives the receivable resource in that offer (*Don’t give him ore, because then he’ll build a city and we’re all doomed*); thus, this is more targeted than a blanket trade embargo,
- 3) Issue a **force-accept** move, that compels the receiver to accept and enact the persuading agent’s trade offer (this is the outcome of successfully persuading the receiver to accept the offer).

The first two moves are of type FOT and the last is of type MT. In our experiments, the strategies for deciding among this expanded set of actions vary.

The JSettlers heuristics already include an FOT type strategy via a trade embargo against players who are believed to be close to winning: i.e., an agent won’t trade with a player who has 8 or more VPs or whose Estimated Time to Win (ETW) is below a set threshold. We disabled this behaviour here in order to avoid any interferences with the issues we are testing. Further, as our experimental results suggest, imposing a trade embargo on the leader when he reaches 8 VPs isn’t effective—it is too late in the game in that the leader is so close to winning that the embargo has insufficient influence on the final outcome.

Since we are testing the upper bound on effectiveness of persuasion moves in various contexts in the game, we start with some unrealistic assumptions about how persuasive an agent can be and then reduce the number of persuasion moves the agent can make and limit the length of time the FOT type of moves are in effect. In this way, we identify the contexts where it is most likely they will have an advantageous effect.

V. EXPERIMENTAL METHOD

A simulation for testing the different persuasion moves consisted of 1 persuading agent playing 3 baseline 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. Running 10,000 games takes about 1h on a current desktop computer.

We investigate persuasion by varying 4 parameters:

- 1) The **type** of move (embargo, block, force-accept),
- 2) The **number** of times an agent can execute that move (modelling how persuasive the agent is in the sense that we vary the number of persuasion arguments he can produce),
- 3) The number of **VPs** the leader has when the persuader starts making persuasion moves (modelling when in the course of a game these moves are most effective),
- 4) The number of **turns** that an embargo or blocking move has an effect (modelling how persuasive the agent is via the stability of the coalition against the embargoed or blocked agent).

In addition to measuring the win rate, we studied the agents’ negotiating and trading behaviour by analysing the total number of trades they made and the number of resources they received by making trades. We will also report the number of persuasion moves they executed. Observe that the number of persuasion moves the agent *actually* makes is a dependent variable on all of the above parameters: this could be fewer than the number he *can* make.

We performed Z-tests to test significance of win rates against the null hypothesis and analysed number of trades and resources received via paired *t*-tests for all combinations of opponents. We used a significance threshold of $p < 0.01$ (when reporting a result as *significant* we will omit the *p*-value). This roughly means that win rates between 0.24 and 0.26 do not differ significantly from the null-hypothesis, which is why we highlight these values in the graphs below.

We report the average numbers for the persuading agent in the 10,000 simulations, and the average numbers across all three baseline agents in 10,000 simulations. 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 instances of the baseline agent, i.e. all differences result from agent modifications.

We use as our baseline agent the one that’s called *ranking* in [7], because (i) it substantially improves over the agent that comes with the *JSettlers* system and (ii) it provides an easier and cleaner implementation of the game-play strategy.

VI. FEWER OPPONENT TRADES I – TRADE EMBARGOES

We start with the persuading agent persuading two of his opponents to enact a trade embargo against the third opponent, and we vary the parameters given above. As already mentioned, the two opponents always comply with this persuasion move. In the simulations, we restrict ourselves to games where at most one trade embargo can be in effect at any given time. This not only avoids the added complexity of deciding on several embargoes simultaneously, but also having only one potential trading partner can hurt the persuading agent (recall the correlation between trading and winning).

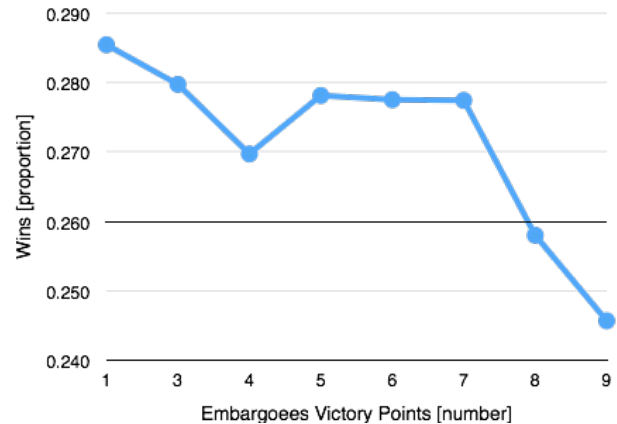


Fig. 2. Effect of how far the game has progressed when an embargo against the leading baseline player is enacted on the win rate of the persuading agent.

We are looking for the point in the game when it is most advantageous for a player to identify a player to embargo and initiate that embargo. Intuitively, blocking the leader when he is already close to winning may be ineffective because he is still likely to win even if he can’t trade. On the other hand, blocking a player early on has a higher chance of targeting someone who is not a contender for winning anyway.

A. Issuing a permanent embargo against one player

The simplest embargo strategy is that the persuading agent issues just one embargo against one opponent and the embargo lasts until the end of the game. The simulations showed that such embargoes are more effective the earlier in the game the persuading agent can issue them, cf. Figure 2.

Issuing an embargo against a player with only 1 VP essentially amounts to randomly choosing an embargoee at the start of the game. So, effectively, the game is not played between 1 persuading and 3 baseline agents but rather between 1 persuading agent, 2 baseline agents and 1 *nontrading* agent. As we have shown previously [6], an agent that is not able to trade is severely handicapped in its ability to win games.

While issuing an embargo when the embargoee only has 1 VP may be most beneficial for the persuading agent, he will be unlikely to identify an argument that convinces his opponents to comply. Such arguments would have to speak to the embargoee’s known, and prior, player type, e.g. *Let’s not trade with him in this game—he always builds to block us*.

On the other hand, Figure 2 shows that there is no big difference to win rates between issuing the embargo when a leader (or a ‘leading follower’ in those contexts where the persuading agent is leading) is at 3 VPs vs. 7 VPs, and certainly at 7 VPs one might convince players to impose embargoes purely on the basis that somebody is currently close to winning (i.e., *Don’t trade with him, he’s going to win*). Note we speak of a leader not *the* leader because multiple players can be in the lead with the same number of VPs, in which case the persuading agent chooses which leader to embargo randomly.

Figure 2 also shows that while such persuasion tactics may be more convincing when the leader has 8 or more VPs, they

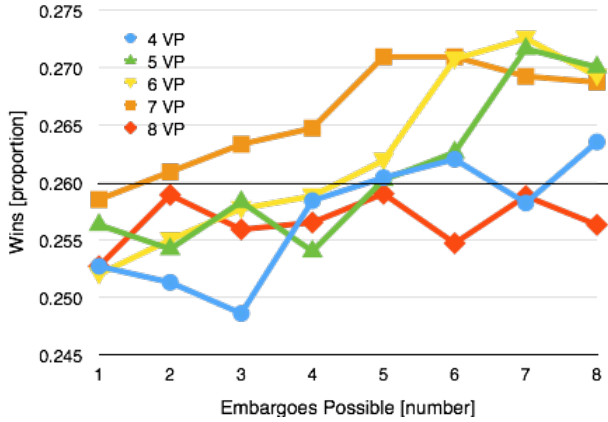


Fig. 3. Effect of the number of embargoes available to the persuading agent. A win rate above 0.26 is a significant benefit.

are now ineffective even if the other players comply: a leader with 8 VPs is likely to win even if he can't trade. So, don't wait too long before attempting this kind of persuasion move! This is notable because the standard behaviour of the *JSettlers* agent is to stop trading with any player who has 8 or more VPs: our experiments expose this to be suboptimal.

As in the simulations reported in [6], the differences in win rate are correlated with the number of trades the agent executes and the number of resources he acquires. Initiating the embargo at 1 VP means that the baseline agents only make 7.44 trades on average compared to the 11.11 of the persuading agent ($t > 46.5$). This results in them receiving only 10.12 resources via trading compared to the 15.05 of the persuader ($t > 45.2$). At 8 VPs these differences are not significant any more (13.93 vs. 14.13 trades, and receiving 18.72 vs. 19.04 resources).

B. Number of embargoes

One of the unrealistic aspects of the above scenario is the unlimited embargo length. Intuitively, the 'coalition' would be less stable, lasting perhaps only one round (i.e., each player has one turn during the embargo period; see the next section for effects of length). Limiting the length enables the persuading agent to initiate more than one embargo in the course of a game: this has the potential advantage that the persuading agent can *change* who is embargoed when there is a change in leading position, thereby adapting to the new context.

Again, we explore the question of when in the game the persuading agent will gain most from initiating these embargoes. For these simulations, the heuristic on who to embargo is slightly more realistic in that only a leader (so not a leading follower when the persuading agent is leading) is ever embargoed—more realistic in that convincing opponents to embargo leaders is easier than convincing them to embargo non-leaders. As before, the persuading agent starts issuing embargoes when the embargoee is a leader with at least the specified number of VPs; if the persuading agent is the leader no embargo is (currently) issued. Embargoes are always issued at the start of the embargoee's own turn.

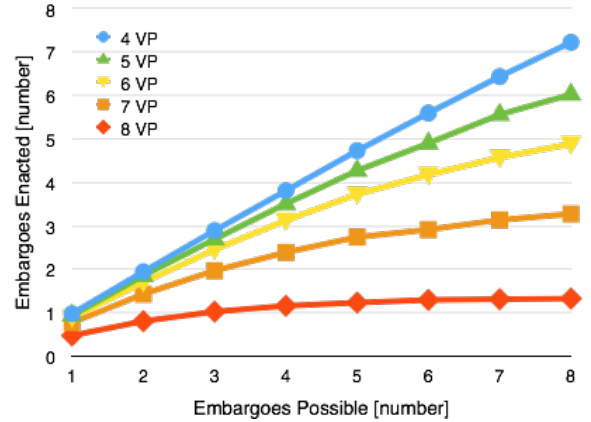


Fig. 4. Number of embargoes enacted depending on how many the persuading agent has available and on when it starts doing so.

Figure 3 shows the results for different numbers of embargoes the persuading agent can issue within a game. To keep these figures easier to read, we omit simulations for fewer than 4 VPs when they do not improve the agent's performance. Not surprisingly, the more embargoes the persuading agent can make, the better his chances of winning. None of the agents significantly improve their win rates if they can issue only 1 embargo, and only the 7 VP agent has a significant advantage with 2 embargoes available (his win rate is 0.261).

This is also the agent that is doing best, when compared to the other persuading types while keeping their number of available embargoes constant. The 5 and 6 VP agents reach his level of performance but only when they have many embargoes available. The 8 VP agent fails to significantly improve his chances of winning games (win rates are all below 0.26).

Two main factors account for these results. The first can be seen by the performance of the 5 and 6 VP agents with 4 available embargoes: the embargoes don't hurt a leader with 5 or 6 VPs enough to help the persuading agent to ultimately win more games. And later on, when a leader has more VPs, the persuading agent has 'used up' his available embargoes! So, these agents issue their embargoes too early. In comparison, the 7 VP agent with the same 4 embargoes available (but issued later) has significantly improved win rates.

The second factor is revealed by the number of embargoes the persuading agents *actually* impose as compared with how many they are allowed to impose, cf. Figure 4: the later in the game the agent starts issuing embargoes, the fewer embargo moves it makes; and the 8 VP agent never enacts more than 1.33 embargoes, even when 8 are allowed. Thus 8 VPs is simply too late in the game: fewer rounds of the game remain than there are available '1-round' embargoes.

Note that the *time* when the embargoes are initiated has a stronger effect on the win rate than the *number* of (actual) embargoes: throughout, the 4 VP agent issues more embargoes than the 7 VP agent, but the 7 VP agent has higher win rates!

With respect to the number of trades and resources received, with 1 embargo available none of the agents significantly improves his performance. With 8 embargoes available, the 4 VP agent increases his number of trades to 13.46 on

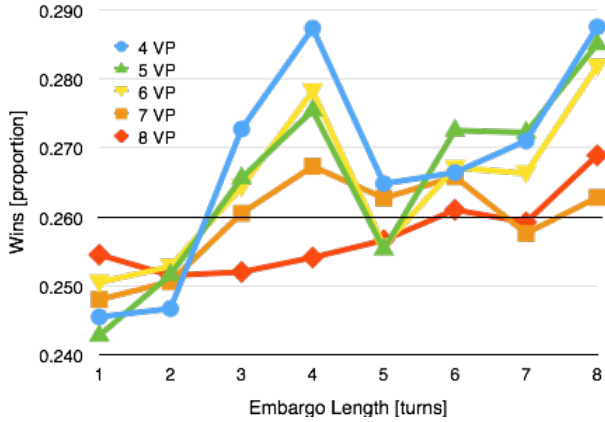


Fig. 5. Effects of embargo length if number of embargoes is unlimited.

average compared to 12.44 of the baseline agent ($t > 14.3$) and his number of received resources to 18.19 compared to 16.78 ($t > 14.0$). Interestingly, the 7 VP agent has *smaller* absolute improvements of the number of trades (13.79 vs. 13.27; $t > 7.1$) as well as of the number of resources received (18.65 vs. 17.84; $t > 8.0$) but still, his win rate is higher! This is further proof that it is the timing of the moves that is crucial.

C. Embargo Length

Up to now, the embargoes lasted either until the end of the game or for 4 turns (i.e., 1 round). Figure 5 confirms our earlier assumption that embargoes are most effective if they last multiples of rounds. If a leader is allowed to trade for 3 turns before the persuading agent considers blocking him again, the persuader benefits considerably less. In these simulations the persuading agent again has an unlimited number of embargo moves at his disposal, but the number of embargoes does not affect the main effect of length, which we confirmed by running two further batches of simulations in which we limited the number of embargoes to 2 and 5.

A 4 VP agent whose embargoes last 1 turn does not achieve significantly more trades or acquire significantly more resources than his baseline opponents. If the 4 VP agent’s embargo length is 8 turns, on the other hand, then the differences are significant: he makes 13.46 trades (vs. 12.44 of the baseline agents; $t > 25.4$) and receives 18.19 resources (vs. 16.78; $t > 25.7$). The differences remain significant even if the agent doesn’t start instigating his 8-turn embargo until the leader has 8 VPs, but only just, with 14.11 vs. 13.91 trades ($t > 2.6$) and 19.03 vs. 18.70 resources received ($t > 2.8$). So, the embargo length directly affects the agents’ chances of winning.

VII. FEWER OPPONENT TRADES II – BLOCKING TRADES

Rather than completely embargoing an agent, e.g. by the mentioned *He’s about to win!*, an agent can also try to stop a particular kind of trade between two opponents, with an argument like: *If you make that trade with him, he’ll build a road and block you.* ‘Natural’ heuristics for identifying trades that a persuading agent might choose to block are:

TABLE I. UPPER BOUNDS FOR DIFFERENT TYPES OF BLOCKING MOVES.

agent	wins	tr. offers	blocks			resourc. by trade
			proposed	followed	trades	
basel.	0.244	23.59	–	4.57	13.00	17.52
heur. 1	0.269	27.77	13.72	–	14.43	19.45
basel.	0.242	26.74	–	4.97	11.92	16.15
heur. 2	0.275	27.66	14.94	–	14.62	19.77
basel.	0.234	22.88	–	10.14	10.55	14.28
heur. 3	0.299	28.60	30.44	–	15.08	20.39

- 1) Trades where the offer is made by an opponent to the persuading agent, but he personally doesn’t want the trade and so decides to block the opponent from getting his desired trade via any other player, too.
- 2) Trades where the offer is *not* made to the persuading agent.
- 3) Both 1 and 2, i.e. the persuading agent blocks all trades between players except those made to him (and perhaps others too) that he wishes to accept. However, if another agent also accepts the offer he may not get the trade.

Blocking moves are made against all players, not just a leader.

The types of situations we are modelling with these blocking moves are ones where the persuading agent wants to prevent an opponent from making a trade that enables a certain action, such as building a particular piece. For this reason, the persuading agent is not just blocking one individual trade—otherwise, the offering agent can gain a similar trade immediately via a slightly modified offer. Rather, the persuading agent blocks any trade in which the agent gets the receivable resources in his trade offer for a specified number of turns (one of our independent variables). For instance, if the trade offer was to receive wood in exchange for clay, then the persuasion move blocks any trade where the offering agent receives wood.

Table I shows the results for when the persuading agent is not restricted by a maximum number of blocking moves or a minimum number of VPs of a leader. All differences between baseline and persuading agents are significant. Roughly speaking, for heuristics 1 and 2 the persuading agent blocks about 20% of the baseline agents’ trade offers (each of the 3 baseline agents makes on average the listed number of trade offers), whereas almost half the trade offers get blocked if the persuading agent adopts heuristic 3. For all heuristics this leads to significant differences in the number of trades and the number of resources the agents manage to obtain via trading.

The persuading agent profits more from heuristic 2 than heuristic 1 but even more if he performs both kinds of moves (heuristic 3). Considering this heuristic mixes two distinct cases (one is a response to a received offer, the other interferes in a negotiation among other players), this is not surprising. As we are establishing upper bounds for the usefulness of these moves, we will only be looking at heuristic 3 in the following.

Figure 6 shows the effectiveness of blocking trades when the persuading agent has an unlimited number of moves at his disposal where the resources are blocked until the end of the game, as well as the more realistic case when he can make 8 blocking moves where the offerer is blocked from asking for

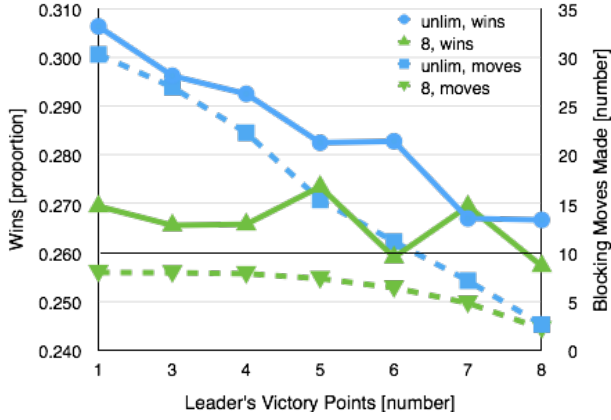


Fig. 6. Effects of blocking moves (using heuristic 3).

the same resources for the duration of 1 turn. We found this length to be most effective and did not find systematic effects when the persuading agent could only make 4 blocking moves or when the resources were blocked for 4 turns.

The results for the more realistic agent having 8 moves lasting 1 turn again show that the best moment to make these moves depends on how far the game has progressed as well as the number of moves available. Considering that an unlimited agent makes ca. 30 blocking moves, being restricted to 8 moves is a severe limitation, for which 6 VPs actually seems to be *not* a good time to start making these moves.

The number of trades and resources received when the agent has an unlimited number of blocking moves available and starts to make these moves from the start of the game are given in Table I. If he starts to make blocking moves at 8 VPs, the advantage shrinks to 14.20 vs. 13.81 trades for the baseline agents ($t > 5.1$) and 19.16 vs. 18.60 resources ($t > 5.0$). If the agent is limited to 8 moves and a block length of 1 turn, at 1 VP he makes 14.39 trades vs. 13.19 ($t > 16$) and receives 19.41 vs. 17.78 resources by trade ($t > 15.5$), which for starting at 8 VPs shrinks to 14.21 vs. 13.84 trades ($t > 5.3$) yielding him 19.15 vs. 18.62 resources ($t > 5.3$).

VIII. MORE TRADES: FORCE-ACCEPT MOVES

Rather than preventing other players from getting resources, a player also benefits from getting more of the trades he wants. We modelled this with *force-accept* moves that compel an opponent to make the trade provided it is possible (i.e., he has the necessary resources to execute the trade).

Figure 7 shows how the win rate of the persuading agent changes depending on the number of force-accept moves he has available and the point in time when he makes them. Clearly, the more force-accept moves, the better! Like embargoes, these moves are most effective when a leader has around 7 VPs. But in contrast to embargoes, they do not depend so much on the stage of the game: they provide significant advantages for the 8 VP agent and, overall, the other agents significantly improve their win rates with fewer moves. Furthermore, the improvements in win rate are slightly larger than for trade embargoes as well as for blocking moves.

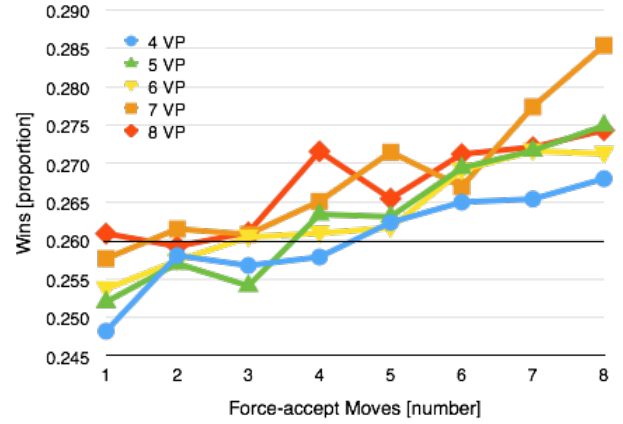


Fig. 7. Effects of force-accept moves.

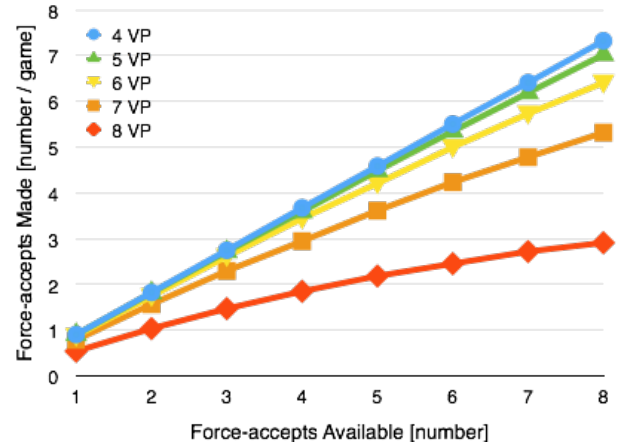


Fig. 8. Number of force-accept moves actually made, depending on how many moves he can make.

Just like embargoes and blocking moves, the later in the game the persuading agent starts making force-accept moves, the fewer such moves he makes; see Figure 8. Also like the other move types, force-accept moves are most efficient (in terms of boosting one's chances to win) for the 7 VP agent.

Looking at the persuading agents' biggest improvements, i.e. when they have 8 moves available, the 4 VP agent makes 18.04 trades (vs. 15.76 for the baseline agent; $t > 27.8$) and receives 22.80 resources by trade (vs. 21.48; $t > 11.6$). Again, the 7 VP agent achieves a bigger improvement in win rate with a slightly smaller increase in trades (17.54 vs. 15.45; $t > 23.9$) and more resources received (22.43 vs. 21.00; $t > 10.7$). Thus once again, the timing is crucial.

If the agent makes a force-accept move without considering whether the trade is actually *possible* (e.g., does the trade partner have the required resources?), the trade may not happen, and he wastes one of the moves at his disposal. In [6] we showed that agents with an accurate belief model of what resources their opponents have are more successful. We, therefore made these moves more targeted: the agent only makes a (valuable) force-accept move if he believes that the other agent has the resources to make the trade.

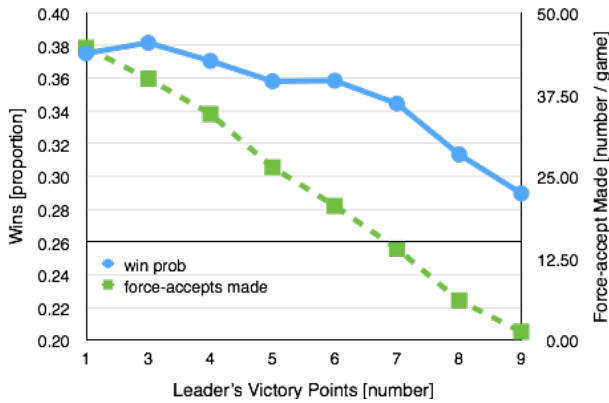


Fig. 9. Win rate and number of force-accept moves made. The number of trade embargoes enacted is always between 0.99 and 1.0.

While this doesn't affect the win rate, the persuading agent uses slightly fewer persuasion moves with the biggest difference being 0.40 for the 7VP agent with 8 available moves. (As these agents do not play against each other directly, we cannot present any statistics.) This further supports what we already found and reported in [6], namely that the agents' standard belief model is quite accurate.

IX. A SUPER-PERSUADER: EMBARGOES & FORCE-ACCEPT

As we are looking for upper bounds of the usefulness of persuasion moves, we gave the persuading agent the ability to make embargo *and* force-accept moves; see Figure 9. This agent achieves a win rate of more than 0.38. Thus, despite all the advantages of persuading, the non-persuading agents still win about 60% of the games! Together with the results we presented in [6], where an agent that didn't trade at all had a win rate of 0.127 and one not making any trade offers (but accepting other players' offers) had a win rate of 0.162, this again shows that the quality of the negotiation and trading strategies in *Settlers* can only increase or decrease an agent's win rate up to a certain point: much of the game is due to chance and agents can always trade with the bank or a port.

The 1 VP agent manages to conclude 35.23 vs. 16.22 trades ($t > 84.2$) to receive 38.63 vs. 22.44 resources ($t > 58.9$); and even the 9 VP agent gets 12.57 vs. 9.10 trades ($t > 41.3$) and 16.70 vs. 12.36 resources ($t > 38.8$).

X. CONCLUSION

In this paper, we showed in a proof by demonstration that one can rapidly obtain a decent symbolic model of persuasion in a complex game, where a player's own preferences over the outcomes of negotiation are incomplete, uncertain, and dynamic. Success critically depends on empirical data via game simulations, which in turn guides the design and development of heuristic strategies that improve the player's chances of winning. Here, we focussed on a specific benefit/cost ratio of persuasion tactics: at what stage in the (complex) game would successfully persuading an agent impact one's chances to win the most? This is a critical question in scenarios where the cognitive effort of identifying a successful persuasion move

itself carries risks (e.g., the risk of 'losing out' while making such calculations to another player that changes the course of the negotiation).

We showed how this empirical approach works in the domain of *Settlers*. In addition, some specific lessons can be learned for this particular domain, some of which are far from obvious *a priori*. For instance, our results reveal a tipping point in the game when the leading player has 7 VPs: at this point it's relatively easy to change the course of the game, but once this is achieved it is increasingly hard for opponents to change it further! The results also suggest that MT type moves benefit the persuading agent slightly more than FOT type moves.

Next, we will empirically investigate persuasion strategies in complex games in contexts where receivers are less gullible, i.e., which persuasion move is most likely to convince a player of a particular type to comply and so change his behaviour?

ACKNOWLEDGMENT

This work is supported by ERC grant 269427 (STAC).

REFERENCES

- [1] N. Asher and A. Lascarides, "Strategic conversation," *Semantics & Pragmatics*, vol. 6, no. 2, pp. 1–62, 2013.
- [2] K. Binmore, *Game Theory and the Social Contract: Just Playing*. MIT Press, 1998.
- [3] A. Rubinstein, "What is the proper role of game theory to other disciplines?" in *Game Theory: 5 Questions*, V. F. Hendricks and P. G. Hansen, Eds. Automatic Press, 2007.
- [4] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations*. Cambridge UP, 2009.
- [5] A. Cadilhac, N. Asher, A. Lascarides, and F. Benamara, "Preference change," 2013, submitted.
- [6] M. Guhe and A. Lascarides, "Effects of belief and memory on strategic negotiation," in *Proceedings of SEMDIAL 2013*, 2013, pp. 82–91.
- [7] —, "Game strategies for The Settlers of Catan," in *Proceedings of CIG 2014*, 2014.
- [8] 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 SEMDIAL 2012*, Paris, 2012.
- [9] M. Guhe and A. Lascarides, "Persuasion in complex games," Submitted, 2014.
- [10] S. Brams, *Negotiation Games: Applying game theory to bargaining and arbitration*. Routledge, 2003.
- [11] C. Camerer, *Behavioural Game Theory: Experiments in Strategic Interaction*. Princeton University Press, 2003.
- [12] A. Rubinstein and J. Glazer, "A study in the pragmatics of persuasion: a game theoretical approach," *Theoretical Economics*, vol. 1, no. 4, pp. 395–410, 2006.
- [13] I. Szita, G. Chaslot, and P. Spronck, "Monte-carlo tree search in Settlers of Catan," in *Advances in Computer Games*, H. van den Herik and P. Spronck, Eds. Springer, 2010, pp. 21–32.
- [14] G. Roelofs, "Monte Carlo Tree Search in a modern board game framework," 2012, research paper available at umimaas.nl.
- [15] M. Pfeiffer, "Machine learning applications in computer games," Master's thesis, Technical University of Graz, 2003.
- [16] K. Teuber, *Die Siedler von Catan: Regelheft*. Stuttgart, Germany: Kosmos Verlag, 1995.
- [17] R. S. Thomas, "Real-time decision making for adversarial environments using a plan-based heuristic," Ph.D. dissertation, Northwestern U., 2003.