OPTIMAL CONTROL IN OPINION DYNAMICS: CONCEPTS, STRUCTURES, PITFALLS, ALGORITHMS

TORSTEN EYMANN, RAINER HEGSELMANN, STEFAN KÖNIG, SASCHA KURZ, CHRISTOPH NIEMANN, AND JÖRG RAMBAU

Abstract.

1. Introduction

(Rainer & alle)

1.1. Related Work.

1.2. Our contribution.

1.3. Outline of the paper.

2. Modeling the Dynamics of Opinions and Their Control

We are interested in opinions that can be represented by mathematical entities. There may be many ways to transform opinions about every conceivable topic into more or less complicated mathematical structures.

In this paper, we will restrict ourselves to the arguably simplest case where an opinion of an individual $i \in I$ can be represented by a real number x_i in the unit interval $[0, 1]$. This can be an opinion about a real quantity measured in percent of a reasonably restricted range (e.g., where is the economic growth of a country next year between $-2\,\%$ and $6\,\%$), or it may be an opinion about the range between extreme positions (e.g., where should be the balance between free markets and public control).

The opinions may change over time subject to a certain *system dynamics*: We assume that time is discretized into *stages* $T := \{0, 1, 2, \dots, N\}$. The opinion of individual $i \in I$ in stage $t \in T$ is denoted by x_i^t. We call, as usual, the vector $\mathbf{x}^t := (x_i^t)_{i \in I}$ the *state* of the system in Stage t. The *system dynamics* $\mathbf{f}^t$ is a vector valued function that computes the state of the system $\mathbf{x}^{t+1}$ as $\mathbf{x}^{t+1} := \mathbf{f}^t(\mathbf{x}^t)$.

Depending on how $\mathbf{f}^t$ is defined, we obtain different models of opinion dynamics. In this paper, we will only consider so-called *stationary* models, where $\mathbf{f}^t$ does not depend on the Stage t. Therefore, from now on, we will drop the superscript t from the notation and write $\mathbf{f}$ for the system dynamics.

2.1. The Average Model.

The motivation for this model is that each individual is in contact with each other in every stage, and each opinion is influenced by each other (including itself) by the same amount. The mathematical model for this is to define $\mathbf{f}$ as the arithmetic mean of all opinions.

This is boring because after the first stage all opinions are equal (consensus).

2.2. **The Lehrer-Wagner Model.** In this model, each individual is again in contact with each other in every stage. The strengths of the influences of opinions on other opinions are given by strictly positive weights w_{ij} with $\sum_{j \in I} w_{ij} = 1$ for all $i \in I$, with the meaning that the opinion of individual i is influenced by the opinion of individual j with weight w_{ij}. The mathematical formulation of this is to define $\mathbf{f} = (f_i)_{i \in I}$ as a weighted arithmetic mean in the following way:

$$f_i(x_1, \ldots, x_n) := \sum_{j \in I} w_{ij} x_j \tag{1}$$

It can be shown that this, in the limit, leads to consensus as well.[1] It leads, as we will see below, still to an interesting optimal control problem.

2.3. **The Bounded-Confidence Model.** The motivation for this model is that our individuals ignore too distant opinions of others. Formally, we fix once and for all an $\epsilon \in (0, 1)$, and each individual is influenced only by opinions that are no more than ϵ away from his or her own opinion. We call $[x_i^t - \epsilon, x_i^t + \epsilon]$ the *confidence interval of individual* i *in Stage* t. Let the *confidence set* $I_i(x_1, \ldots, x_n)$ of individual $i \in I$ in state $\mathbf{x} = (x_1, \ldots, x_n)$ be defined as

$$I_i(x_1, \ldots, x_n) := \{j \in I : |x_j - x_i| \leqslant \epsilon\}. \tag{2}$$

Then the system dynamics of the bounded confidence model is given as follows:

$$f_i(x_1, \ldots, x_n) := \frac{1}{|I_i(x_1, \ldots, x_n)|} \sum_{j \in I_i(x_1, \ldots, x_n)} x_j. \tag{3}$$

This system dynamics is mathematically extremely interesting because it is not even continuous. It can be argued about whether it makes sense that confidence is dropped completely at a very sharp point in opinion space or whether a more continuous transition should be chosen. Our personal experience is that it actually happens that all of a sudden, because of some event, we lose our faith in some person; this supports the non-continuous characteristics of the bounded-confidence model. A possible extension might be a stochastic disturbance on ϵ, but, as we will see, bounded-confidence is still far from being completely understood. Therefore, in this paper bounded confidence will be in the main focus.

2.4. **A New Opinion Control Model.** Given a dynamical system as above, we can of course think about the possibility of a control, that can influence the system dynamics. Formally, this means that the system dynamics $\mathbf{f}$ depends also on some additional exogenously provided data $\mathbf{u}$, the *control*.

The easiest way to think about opinion control is to carefully state an opinion in front of all individuals so that the new opinion takes part in influencing all the individuals opinion. Formally, the controller can place one or more additional opinions in the opinion space in order to guide the individuals opinions in a specified direction. One possible interpretation of this to place suitable statements in the stages of a marketing campaign in order to convince as many customers as possible to buy the product rather than the competition; another is to present political speeches with consciously designed opinions during the stages of an election campaign in order to get as many votes as possible, i. e., the opinions of as many as possible voters are closer to the party than to the competition.

Formally, this means in the simplest case (and we will restrict to this case) that the controller can present an additional opinion u^t in every stage that takes part in the opinion dynamics. The corresponding system dynamics, taking the control as an additional argument, are then given as follows (with $x_0 := u$ and $I_0 := I \cup \{0\}$ as well as w_{ij} this time with $\sum_{j \in I_0} w_{ij} = 1$ for easier notation):

[1] This is an easy consequence of the Banach Fixed Point Theorem, since this dynamics is a contraction.

$$f_i(x_0; x_1, \dots, x_n) := \sum_{j \in I_0} w_{ij} x_j \qquad \text{(Lehrer-Wagner-Control)}$$

$$f_i(x_0; x_1, \dots, x_n) := \frac{1}{|I_i(x_0, x_1, \dots, x_n)|} \sum_{j \in I_i(x_0, x_1, \dots, x_n)} x_j. \qquad \text{(Bounded-Confidence-Control)}$$

We can interpret this as a usual model of opinion dynamics with an additional opinion x_0 that can be positioned freely in every stage by the controller.

We still have to formalize what the goal of the controller is. In marketing or politics it would be desirable that after a certain number of stages as many opinions as possible are closer to a target opinion (favor the product over the competition, vote for the party) than to competing opinions. In the case of fixed target opinions and fixed competing opinions this can always be expressed as follows: Control opinions in a way such that after N stages there are as many opinions as possible in a given interval$[\ell, r] \subseteq [0, 1]$.

To formalize this, fix an interval $[\ell, r]$ (the *conviction interval*), and let *the conviction set* $J(x_1, \dots, x_n)$ denote the set of all individuals $j \in I$ with $x_j \in [\ell, r]$. We want to maximize the number of convinced individuals. Thus, the problem we want to investigate is the following deterministic discrete-time optimal control problem:

$$\max_{x_0^0, x_0^1, \dots, x_0^{N-1}} |J(x_1^N, \dots, x_n^N)|$$

subject to

$$x_i^{t+1} = f_i(x_0^t; x_1^t, \dots, x_n^t) \qquad \forall t = 0, 1, \dots, N-1, \qquad \text{(System Dynamics)}$$

$$x_0^t \in [0, 1] \qquad \forall t = 0, 1, \dots, N-1, \qquad \text{(Control Restrictions)}$$

where $\mathbf{f} = (f_i)_{i \in I}$ is one of the controlled system dynamics in Equations (Lehrer-Wagner-Control) and (Bounded-Confidence-Control), resp.

2.5. **Our Benchmark Example.** We now design a special instance of our optimal control problem that serves as our benchmark for computational investigations. We are given eleven individuals with starting opinions $0, 0.1, 0.2 \dots, 0.9, 1$. Our conviction interval is the interval $[0.375, 0.625]$. The goal is to maximize the number of convinced individuals in stage $1, 2, \dots, 10$, resp.

We will see, that even for this innocent-looking example we were not able to find the optimal number of convinced individuals for all numbers of stages between 1 and 10.

3. Simulation and Pitfalls from Numerical Mathematics

In this section we want to convince the reader from the fact that the numerical inaccuracies even in a simple simulation of the bounded confidence model (as opposed to the Lehrer-Wagner model) have drastic effects on the results observed. The only way for us to cope with this problem is to resort to exact rational arithmetic's throughout, although there may be more sophisticated methods to improve efficiency. This numerical instability has the more serious consequence that of-the-shelf optimization algorithms with floating point arithmetic's can not be used without checking the results for correctness in exact arithmetics.

Let us start with an example of 6 individuals with opinions being regularly distributed at the positions 0.0, 0.2, $\dots$, 1.0 with $\epsilon = 0.2$. For a moment we forget on controlling the opinion dynamics and focus on the consensus process. Since there is a mirror symmetry to opinion 0.5 and we assume no external control that might destroy this symmetry there should be such a symmetry in each stage.

Using exact arithmetic we obtain that the opinions of our individuals are given by

$$
\begin{aligned}
x^0 &= (0.0, 0.2, 0.4, 0.6, 0.8, 1.0),\\
x^1 &= (0.1, 0.2, 0.4, 0.6, 0.8, 0.9),\\
x^2 &= (0.15, 0.2\overline{3}, 0.4, 0.6, 0.7\overline{6}, 0.85),\\
x^3 &= (0.191\overline{6}, 0.26\overline{1}, 0.4\overline{1}, 0.5\overline{8}, 0.73\overline{8}, 0.808\overline{3})\\
&= \left(\frac{23}{120}, \frac{47}{180}, \frac{37}{90}, \frac{53}{90}, \frac{133}{180}, \frac{97}{120}\right),\\
x^4 &= \left(\frac{163}{720}, \frac{311}{1080}, \frac{227}{540}, \frac{313}{540}, \frac{769}{1080}, \frac{557}{720}\right),\\
x^5 &= \left(\frac{673}{2160}, \frac{673}{2160}, \frac{3271}{8640}, \frac{5369}{8640}, \frac{1487}{2160}, \frac{1487}{2160}\right),\\
x^6 &= \left(\frac{577}{1728}, \frac{577}{1728}, \frac{577}{1728}, \frac{1151}{1728}, \frac{1151}{1728}, \frac{1151}{1728}\right)\\
&= \Big(0.333912\overline{037}, 0.333912\overline{037}, 0.333912\overline{037}, 0.666087\overline{962},\\
&\quad 0.666087\overline{962}, 0.666087\overline{962}\Big),
\end{aligned}
$$

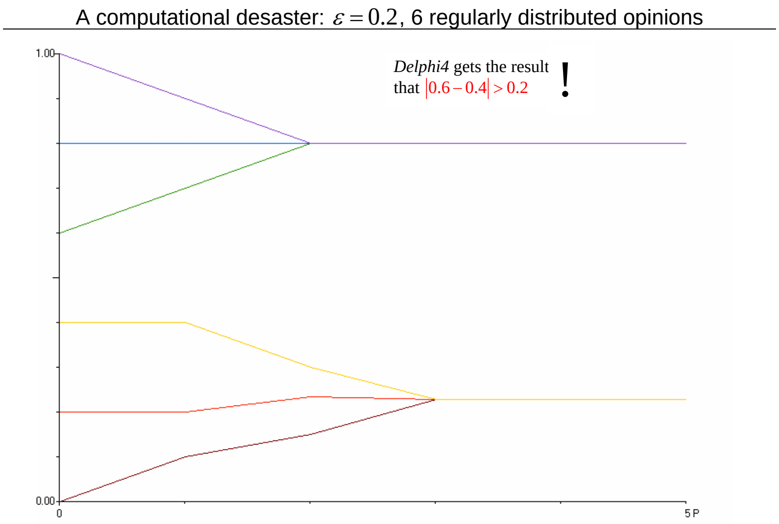

FIGURE 1. A computational desaster.

With the computation at hand we can see that this simple example leads to a quite complex dynamic outlasting 6 stages before ending in a consensus of two different opinions. In Figure 1 we have depicted what a Delphi programm makes out of this tiny innocent looking example.

We observe several things. The plot of Figure 1 seems not to correlate with the exact numbers in any sense. The Delphi program determines two different opinions in the consensus which is attained after only three stages. There is another interesting detail hidden in the graphics. At the starting stage individual 3 is influenced by individuals 2, 3, and 4, but individual 4 is influenced by individuals 4 and 5. So the influence is asymmetric. Individual 4 influence individual 3 but individual 3 does not influence individual 4. Clearly such an inconsistency could easily be avoided by determining influence variables $c_{i,j}$ which are equal to 1 is i influences j and 0 otherwise. The symmetry could be forced by setting $c_{ij} = c_{ji}$. So the last observation is a nasty pitfall one might step into, which could on the other hand be easily avoided. You may ask yourself how you would implement the described dynamics in your favorite programming language. Just for fun you may check whether $|x_i - x_j| \leqslant \epsilon$ always gives the same result as $|x_j - x_i| \leqslant \epsilon$ if you implement the dynamics more directly.

On the other hand the definition of the bounded confidence model requests hard decision as $|x_i - x_j| \leqslant \epsilon$ or $|x_i - x_j| > \epsilon$. It is very common to use floating point numbers in those simulations. If the functional equations behave very continuous nothing is wrong about that. Unfortunately in our situation the dynamics is very discontinuous. A small jiggle on one of the intermediate results may end in such drastic effects as Figure 1 compared to the *real* result. The only patch that came to our mind which was capable of dealing with the numerical instability was to use exact arithmetic. This means that we represent all numbers as fractions where the numerator and the denominator are integers with unlimited accuracy. We remark that we have used the Class Library of Numbers (CLN) a C++-package, but similar packages should be available also for your favorite programming languages.

In the starting phase of our investigation in optimal control of opinion dynamics we have also used floating point arithmetic for a short time period. By some spontaneous guesses we came up with the control

$$[0.35, 0.3875, 0.775, 0.439583, 0.656771, 0.618083, 0.588083, 0.558083, 0.550413, 0.504684].$$

Our computer program tells us that we would achieve 10 convinced individuals by applying this control vector on the benchmark problem of Subsection 2.5. This would have been rather good, since i. e. a genetic algorithm achieves only 9 convinced individuals, see Subsection 6.2. Unfortunately a recalculation using exact arithmetic yields only 4 convinced individuals, which is a rather bad control. We have made similar experiences with all considered heuristics or algorithms being described in this article.

There a quite a lot of articles dealing with the simulation of the bounded confidence model. To our knowledge none of these mentioned the use of exact arithmetic. So one could assume that the authors have used ordinary floating point numbers with limited precision for there considerations. It is an interesting question whether all of these obtained results remain more or less the same if being recalculated with exact arithmetic.

- Examples (with graphics) for numerical instability (Rainer,Sascha)
- Comparison with Lehrer-Wagner (Rainer)

4. Basic Structural Properties of Optimal Controls

In this section we collect some basic facts about structural properties of optimal controls. First of all we mention that one or usually a whole set of optimal controls do exist. The number of convinced individuals is in any stage trivially bounded from above by the total number of individuals $|I|$. So to every control there corresponds a bounded integer valued number of convinced individuals.

The next thing we observe is that given enough time (number of stages) we could always achieve this upper bound of $|I|$ convinced individuals. Therefore we may determine the index i

of the individual with the leftmost opinion and place the control at $x_i + \epsilon$ in every stage. After a finite number of stages all individuals have the same opinion. After this preparing period we can move the whole set of individuals I by an amount between 0 and $\frac{\varepsilon}{|I|+1}$ in every direction we like. So again after a finite number of stages all individuals are in the conviction interval. After that time point we could simply place the control x_0 directly on the opinion x_i of all individuals. We remark that one might consider some subtlety which may occur if $x_i + \epsilon > 1$. In this case we simply set $x_0 = 1$. An easy estimation yields that after at most $2 + \frac{2}{(|I|+1)\cdot\epsilon}$ stages all $|I|$ individuals have an opinion in the conviction interval. By setting $[l, r] = [0, 0]$ and $x_i^0 = 1$ for all $i \in I$ we easily see that this estimation gives the right order of magnitude.

The opinion control in the bounded confidence model, as introduced here, is somewhat to hard to be analyzed analytically at our current state of knowledge. Therefore we introduce a slightly modification of the problem which seems to be a bit easier. So far we had a conviction interval $[l, r]$. Now we restrict ourselves to intervals of the form $[l, 1]$ or $[0, r]$. So the problem becomes something like a rope-pulling game with center c. Let assume that we have the very special case where all individuals have the same opinion and let us denote the distance $|x_i - c|$ between the center of the rope-pulling game and the opinion of the individuals by δ. If there are $r \geqslant \frac{\delta}{(|I|+1)\cdot\epsilon}$ stages left, then all $|I|$ individuals can be convinced by an optimal control, otherwise no individual can be convinced after r stages. This fact can be seen easily as follows. Whatever our vector of controls is, all pairs of individuals (i, j) will have equal opinions $x_i^t = x_j^t$ for ever stage $t \in \mathbb{N}$. By suitably placing the control we can move the opinion of all individuals by an amount between 0 and $\frac{\varepsilon}{|I|+1}$ in every direction we like. Thus we have

$$x_i^t \in \left[\max\left(x_i^0 - t \cdot \frac{\varepsilon}{|I|+1}, 0\right), \min\left(x_i^0 + t \cdot \frac{\varepsilon}{|I|+1}, 1\right)\right]$$

and each value of this interval could be achieved by x_i^t. For the more general case of an arbitrary conviction interval we have only to check if it intersects with the stated interval.

We remark that it is possible, in principle, to analyze the situation for at least $|I| = 2$ completely. If we manage to present it in a readable way this will be part of a forthcoming article.

5. A Mathematical Model for the Optimal Control Problem

In this section we describe a mathematical model, namely a mixed integer linear programming model (MILP), for our optimal control problem. This serves two purposes:

- By reducing the number of periods in the model, we can exactly solve the model. This allows us to employ a receding-horizon heuristic to the original problem.
- The model allows for the computation of performance bounds, in our case upper bounds on the number of supporters that can be gathered in the final state by any control.

The model is — not surprisingly — much more powerful in the Lehrer-Wagner model than in the bounded-confidence model; the former takes profit of the linear system dynamics whereas the latter suffers a lot from the highly non-continuous system dynamics and the numerical instability. More specifically, our model is not able to represent the original problem exactly. We will provide, however, actually two models: one is correct in the sense that every upper bound on the objective value of the model is an upper bound on the optimal number of convinced individuals in the original problem (but possibly not vice versa); the other one is correct in the sense that any feasible solution to it is a feasible solution to the original problem (but possibly not vice versa).

Why do we choose a MILP formulation as a model? Our original problem does not look like a mixed integer problem at all. Well, MILP is at least able to capture combinatorial structures like the confidence set or the conviction set, and the objective value of mixed integer programs is in general not continuous in the input data. After all, frankly speaking, we had no better idea.

The motivation for using integral variables in a model for our optimal control problem is that the dynamics mainly depends on the structures of the confidence sets and the conviction sets: We can use binary variables to indicate how the conviction sets and the confidence sets, resp., look like.

Since the bounded confidence model requires some experience in modeling with MILPs, we start with a model for the Lehrer-Wagner optimal control problem. Later on, when the main principles are explained, we will present a model for the bounded-confidence case.

5.1. **MILP Model for the Lehrer-Wagner Optimal-Control Problem.** The following simple model is quite a standard strategy in MILP. We first list the *variables* of the model.

- The continuous variables $x_0^t \in [0,1]$, $t = 0,1,\dots,N-1$ denote the positions in opinion space where we place a control in the various stages; these are the variables that we are really after.
- The continuous variables $x_i^t \in [0,1]$, $i \in I$, $t = 0,1,\dots,N$ denote the positions of the individuals in the various stages; these variables measure the system states. The variables in Stage 0 are given as input data (start state).
- For each individual, we want to measure whether its position in Stage N is inside the conviction interval; to this end, we use binary variables $z_i \in \{0,1\}$, $i \in I$, with the following meaning: $z_i = 1$ if and only if i is convinced in Stage N, i.e., $x_i^N \in [\ell, r]$.

With this, we may formulate the *goal* of the model: we want to maximize the number of convinced individuals, which can be expressed as follows:

$$\max \sum_{i \in I} z_i \tag{4}$$

Now, the success measuring variables z_i have to be coupled with our decisions x_0^t via the system states and the system dynamics. The following linear *side constraint* couples the decisions to the system states:

$$x_i^{t+1} = \sum_{i \in I_0} w_{ij} x_j^t \quad \text{for all } i \in I, t = 0,1,\dots,N-1. \tag{5}$$

So far, we did not restrict the binary variables. A solver would simply set the all to 1 and achieve an objective value of n (all convinced), because the binary variables so far have nothing to do with the underlying dynamical system.

The binary variables can now be coupled to the system state variables in Stage N by a standard MILP modeling trick as follows. The logical implication must be: If $z_i = 1$, i.e., if we want to count an individual as convinced, then $\ell \leqslant x_i^N \leqslant r$ must hold. In other words, the inequalities $\ell \leqslant x_i^N \leqslant r$ can be violated when $z_i = 0$, but they must be satisfied whenever $z_i = 1$.

We show the trick for the inequality $\ell \leqslant x_i^N$, the other inequality can be handled analogously. The maximal violation of the inequality $\ell - x_i^N \leqslant 0$ is ℓ, since $\ell - x \leqslant \ell$ for all $x \in [0,1]$. That means, the inequality $\ell - x_i^N \leqslant \ell$ does trivially hold, no matter where x_i^N is in $[0,1]$. We want to impose the trivial inequality $\ell - x_i^N \leqslant \ell$ whenever $z_i = 0$ and the non-trivial inequality $\ell - x_i^N \leqslant 0$ whenever $z_i = 1$. But this can be achieved in one step by imposing the inequality

$$\ell - x_i^N \leqslant \ell(1 - z_i) \quad \text{for all } i \in I. \tag{6}$$

The analogously derived inequality for the right border of the conviction interval reads

$$x_i^N - r \leqslant (1 - r)(1 - z_i) \quad \text{for all } i \in I. \tag{7}$$

The complete MILP reads as follows:

$$\max \sum_{i \in I} z_i \tag{8a}$$

subject to

$$x_i^{t+1} = \sum_{i \in I_0} w_{ij} x_j^t \qquad \text{for all } i \in I, t = 0, 1, \dots, N-1 \tag{8b}$$

$$\ell - x_i^N \leqslant \ell(1 - z_i) \qquad \text{for all } i \in I \tag{8c}$$

$$x_i^N - r \leqslant (1 - r)(1 - z_i) \qquad \text{for all } i \in I. \tag{8d}$$

These types of MILPs can be solved quite efficiently by of-the-shelf software like ILOG cplex. In particular, solving our benchmark problem for any number of stages between 1 and 10 is absolutely no problem. We will present computational results later in this paper.

5.2. **MILP Model for the Bounded-Confidence Optimal-Control Problem.** A MILP for the Bounded-Confidence dynamics can be designed by essentially applying the modeling trick for counting the convinced individuals over and over again. This leads to a model that, in our experience, could not even be solved by ILOG cplex for our benchmark problem for $N = 3$.

We have found a MILP with a larger number of variables that could be solved by ILOG cplex for our benchmark problem up to $N = 5$. We will present this model in the following without going into too much detail.

For our model, we assume that all individuals are numbered according to their starting opinion, i.e., $i < j$ implies $x_i^0 \leqslant x_j^0$ for $i, j \in I$.

Observation 5.1. *The order of individuals in opinion space does not change, i.e.,f* $x_i^t \leqslant x_j^t$ *for some* $i, j \in I$ *and some* $t = 0, 1, \dots, N-1$*, then* $x_i^{t+1} \leqslant x_j^{t+1}$.

The model uses the following variables:

- The control variables x_0^t, $t = 0, 1, \dots, N-1$ are as above.
- Similarly, the state variables x_i^t, $i \in I$, $t = 0, 1, \dots, N$ are as above.
- For $j_{\min}, j_{\max} \in I$ and $l, r \in \{0, 1\}$, we introduce variables $v^t_{i,(j_{\min},j_{\max};l,r)}$ where $v^t_{i,(j_{\min},j_{\max};l,r)} = 1$ if and only if the following holds: Individual $j_{\min}$ is the minimal index of an individual in the confidence interval of i, Individual $j_{\max}$ is the maximal index of an individual in the confidence interval of i, Index $l = 1$ if and only if $x_0^t \geqslant x_i^t - \epsilon$ (i.e., the control is not to the left of the confidence interval of individual i), and Index $r = 1$ if and only if $x_0^t \leqslant x_i^t + \epsilon$ (i.e., the control is not to the right of the confidence interval of individual i). In particular, all variables $v^t_{i,(j_{\min},j_{\max};0,0)}$ must be zero. The motivation for these variables is that they are indicating the unique *combinatorial confidence configuration* $(j_{\min}, j_{\max}; l, r)$ of an individual: If $v^t_{i,(j_{\min},j_{\max};l,r)} = 1$ then we know by Observation 5.1 that all individuals $j \in I$ with $j_{\min} \leqslant j \leqslant j_{\max}$ influence i and that the current control influences i if and only if $l = r = 1$. In MILP language, these variables are *assignment variables* that assign to each individual a unique combinatorial confidence configuration.
- For $j_{\min}, j_{\max} \in I$, we introduce variables $p_{(j_{\min},j_{\max})}$ where $p_{(j_{\min},j_{\max})} = 1$ if and only if the following holds: $j_{\min}$ is the minimal index of an individual in the conviction interval in Stage N, and $j_{\max}$ is the maximal index of an individual in the conviction interval in Stage N. The motivation for these variables is that they are indicating the unique *combinatorial conviction configuration* $(j_{\min}, j_{\max})$ in the final stage: If $p_{(j_{\min},j_{\max})} = 1$ then the number of convinced individuals in Stage N is simply $j_{\max} - j_{\min} + 1$.

We introduce a very small $\delta \geqslant 0$ (in our computational experiments we chose $\delta = 10^{-6}$) with the following meaning: whenever j is not in the confidence interval of i, then $|x_i - x_j| \geqslant \epsilon + \delta$

must hold. This is stronger than the original condition, which is: if j is not in the confidence interval of i, then $|x_i - x_j| > \epsilon$ must hold, and vice versa. This original condition is a strict inequality that can not be handled directly in MILPs, and a transformation to a different MILP (in modified so-called homogeneous variables) is usually numerically highly unstable.

With the modified condition we can choose either to exclude potentially feasible solutions (this happens for $\delta > 0$) or we grant the optimization algorithm to choose freely whether or not j is in the confidence interval of i whenever $|x_i - x_j| = \epsilon$ (this happens for $\delta = 0$). So, whenever we are after feasible solutions we will set δ to something strictly positive, and whenever we are after upper bounds on the optimal number of convinced individuals we set δ to zero.

This slight inaccuracy in modeling is acceptable since the MILP-solvers we can employ use inexact floating point arithmetic with an accuracy of 10^{-6} anyway.

The resulting model can be formulated as follows:

$$\max \sum_{(j_{\min} \leqslant j_{\max})} p_{(j_{\min}, j_{\max})} \tag{9a}$$

subject to

$$\sum_{\substack{j_{\min} \leqslant i \leqslant j_{\max} \\ l,r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; l, r)} = 1 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I \tag{9b}$$

$$x^t_i - x^t_{j_{\min}} - \epsilon - (1 - \epsilon)\left(1 - \sum_{\substack{j_{\max} \geqslant i \\ l,r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; l, r)}\right) \leqslant 0 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I,\ j_{\min} \leqslant i \tag{9c}$$

$$-x^t_i + x^t_{j_{\max}} - \epsilon - (1 - \epsilon)\left(1 - \sum_{\substack{j_{\min} \leqslant i \\ l,r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; l, r)}\right) \leqslant 0 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I,\ j_{\max} \geqslant i \tag{9d}$$

$$x^t_i - x^t_{j_{\min}-1} - \epsilon - \delta + (1 + \epsilon + \delta)\left(1 - \sum_{\substack{j_{\max} \geqslant i \\ l,r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; l, r)}\right) \geqslant 0 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I,\ 0 < j_{\min} \leqslant i \tag{9e}$$

$$-x^t_i + x^t_{j_{\max}+1} - \epsilon - \delta + (1 + \epsilon + \delta)\left(1 - \sum_{\substack{j_{\min} \leqslant i \\ l,r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; l, r)}\right) \geqslant 0 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I,\ i \leqslant j_{\max} < n \tag{9f}$$

$$x^t_i - x^t_0 - \epsilon - (1 - \epsilon)\left(1 - \sum_{\substack{j_{\min} \leqslant i \leqslant j_{\max} \\ r \in \{0,1\}}} v^t_{i,(j_{\min}, j_{\max}; 1, r)}\right) \leqslant 0 \qquad \text{for all } t = 0, 1, \ldots, N-1,\ i \in I \tag{9g}$$

$$-x_i^t + x_0^t - \epsilon - (1-\epsilon)\left(1 - \sum_{\substack{j_{\min}\leqslant i\leqslant j_{\max}\\ l\in\{0,1\}}} v^t_{i,(j_{\min},j_{\max};l,1)}\right) \leqslant 0 \qquad \text{for all } t = 0,1,\dots,N-1,\ i \in I \tag{9h}$$

$$x_i^t - x_0^t - \epsilon - \delta + (1+\epsilon+\delta)\left(1 - \sum_{\substack{j_{\min}\leqslant i\leqslant j_{\max}\\ r\in\{0,1\}}} v^t_{i,(j_{\min},j_{\max};0,r)}\right) \geqslant 0 \qquad \text{for all } t = 0,1,\dots,N-1,\ i \in I \tag{9i}$$

$$-x_i^t + x_0^t - \epsilon - \delta + (1+\epsilon+\delta)\left(1 - \sum_{\substack{j_{\min}\leqslant i\leqslant j_{\max}\\ l\in\{0,1\}}} v^t_{i,(j_{\min},j_{\max};l,0)}\right) \geqslant 0 \qquad \text{for all } t = 0,1,\dots,N-1,\ i \in I \tag{9j}$$

$$\sum_{j_{\min}\leqslant j_{\max}} p_{(j_{\min},j_{\max})} \leqslant 1 \qquad \text{for all } t = 1,\dots,N \tag{9k}$$

$$\ell - x^t_{j_{\min}} - \ell\left(1 - \sum_{j_{\max}\geqslant j_{\min}} p_{(j_{\min},j_{\max})}\right) \leqslant 0 \qquad \text{for all } t = 1,\dots,N \tag{9l}$$

$$x^t_{j_{\max}} - r - (1-r)\left(1 - \sum_{j_{\min}\leqslant j_{\max}} p_{(j_{\min},j_{\max})}\right) \leqslant 0 \qquad \text{for all } t = 1,\dots,N \tag{9m}$$

$$\begin{aligned}(j_{\max} - j_{\min} + 1 + lr)x_i^t - lrx_0^{t-1} - \sum_{j\in I: j_{\min}\leqslant j\leqslant j_{\max}} x_j^{t-1}&\\ - (j_{\max} - j_{\min} + 1 + lr)\big(1 - v^t_{i,(j_{\min},j_{\max};l,r)}\big) \leqslant 0& \qquad \text{for all } t = 1,\dots,N,\ i\in I,\ j_{\min}\leqslant j_{\max},\ l,r\in\{0,1\}\end{aligned} \tag{9n}$$

$$\begin{aligned}(j_{\max} - j_{\min} + 1 + lr)x_i^t - lrx_0^{t-1} - \sum_{j\in I: j_{\min}\leqslant j\leqslant j_{\max}} x_j^{t-1}&\\ + (j_{\max} - j_{\min} + 1 + lr)\big(1 - v^t_{i,(j_{\min},j_{\max};l,r)}\big) \geqslant 0& \qquad \text{for all } t = 1,\dots,N,\ i\in I,\ j_{\min}\leqslant j_{\max},\ l,r\in\{0,1\}\end{aligned} \tag{9o}$$

$$x_i^t \in [0,1] \qquad \text{for all } t = 1,\dots,N,\ i \in I \tag{9p}$$

$$v^t_{i,(j_{\min},j_{\max};l,r)} \in \{0,1\} \qquad \text{for all } t = 1,\dots,N,\ j_{\min}\leqslant i\leqslant j_{\max},\ l,r\in\{0,1\} \tag{9q}$$

$$p_{(j_{\min},j_{\max})} \in \{0,1\} \qquad j_{\min}\leqslant i\leqslant j_{\max} \tag{9r}$$

We note that some perturbation of the objective function based on the average distance of individuals not yet in the conviction interval over all stages helps the MILP solver to find feasible solutions. This problem, however, is much more successfully addressed in the following section.

6. Heuristics to Find Good Controls

In the previous section we have described a MILP-formulation of our problem. So in principle one could solve every problem instance by standard of-the-shelf software like ILOG cplex. In contrast to the Lehrer-Wagner model, where we could solve our benchmark problem for any number of stages between between 1 and 10 without any difficulty, the instances from the bounded confidence model are quite harder. Using the MILP-formulation of the previous section we were only able to determine the optimal control up to 5 stages using ILOG cplex.

The approach using a MILP-formulation has the great advantage that we receive upper bounds for the optimal control. For the other direction we need heuristics that can efficiently determine good control. Also our MILP-approach benefits from good feasible solutions especially if they respect branched variables in the branch & bound search tree. So in the next subsections we give three heuristics to find good controls.

6.1. **Strongest-guy heuristics.** What makes the problem hard is, despite apart from the discontinuous dynamics and numerical instabilities, is the fact, that the control x_0 is a continuous variable in all stages. So at first sight the problem is not a finite one. Using the MILP-approach it becomes a finite problem nevertheless. Let us relax our problem a bit by allowing only a finite number of possibilities for x_0 at any stage and have a closer look at the situation.

By placing a control x_0 at a certain stage some individuals are influenced by x_0 others are not influenced by x_0. We notice that the magnitude of influence rises with the distance between the individuals opinion x_i and the control x_0 as long as their distance remains below ϵ. So the idea is, nevertheless we are not knowing what we are doing, we will do it with full strength. Let c be the center of the conviction interval $[l, r]$ then the set of possible positions of x_0 at a certain stage is given by

$$\begin{cases} x_i + \epsilon & \text{if} \quad x_i \leqslant p, \\ x_i - \epsilon & \text{if} \quad x_i > p \end{cases}$$

for all $i \in I$. We call this relaxation of the problem the strongest-guy heuristics.

Stages	Convinced Individuals	Control
0	3	$[]$
1	3	$[4]$
2	4	$[4, 4]$
3	5	$[4, 8, 3]$
4	6	$[3, 3, 8, 6]$
5	6	$[3, 3, 8, 7, 9]$
6	6	$[3, 7, 3, 8, 3, 3]$
7	8	$[3, 8, 1, 2, 4, 1, 1]$
8	8	$[4, 5, 10, 9, 9, 1, 1, 1]$
9	8	$[3, 9, 3, 8, 9, 8, 8, 8, 9]$
10	11	$[3, 11, 4, 6, 9, 8, 8, 8, 6, 1]$

Table 1. Results of the strongest-guy heuristic on the benchmark example.

Instead of giving the exact values of x_0 for all stages we can also give the indices i if we use this heuristic. In Table 1 we give for our benchmark problem the maximum number of convinced individuals that can be achieved by using the strongest guy heuristic together with the corresponding index-vector. Although being a heuristic, we can conclude that if the number of stages is at least 10 then the maximum number of convinced individuals is 11. Clearly, this is only possible due to the trivial upper bound of $11 = |I|$. The values of Table 1 are proven to be optimal for up to 5 stages by using the MILP approach from the previous section. We remark that we are not aware of any improvements to Table 1.

Having a look at the optimal controls of Table 1 one gets an impression of the hardness of our problem. There seems to be no obvious pattern in the optimal controls. Who would have guessed a control like $[3, 11, 4, 6, 9, 8, 8, 8, 6, 1]$? We remark that starting with x_1 or x_2 at the first stage does not lead to 11 convinced individuals at the end. I.e. starting with x_2 and continuing with x_{11} does lead to only 7 convinced individuals using the strongest-guy heuristic.

We remark that the strongest-guy heuristic can be easily adopted to the situation where some of the $0-1$ variables from the MILP formulation from the previous section are fixed to either 0 or 1.

6.2. **Genetic algorithm.** To get a better idea about the solution space, we implemented a genetic algorithm (GA) to search for optimal solutions. The GA used standard GA methods to evolve good solutions from a randomly chosen set of starting strategies. To use GA for the the problem, the problem has to be formulated to suit the GA terminology. A GA instance consists of *genes* that form a *chromosome*. Each chromosome can be evaluated using a *fitness function*, which serves to determine the chromosome's quality with respect to the original problem. The GA uses alternating steps of *evolution* and *selection* to modify the chromosomes and moving the entire *population* of chromosomes to increase the number of high quality chromosomes. By means of the *survival of the fittest*, the selection process sorts out weak chromosomes with low fitness values, while it retains chromosomes with high fitness values. The remaining chromosomes evolve to the next round of the GA. The following subsection explain the different parts of the GA individually.

6.2.1. *Set Up of the GA.* The only free variables in our example problem are the ten different positions of our freely selectable individual: The ten different positions make up one strategy to control the remaining individuals. In the GA, one strategy is encoded as one chromosome with the ten different positions each occupying one gene. Because of the numerical inaccuracies (Section 3) in standard floating point implementations, the GA had to use fractions to specify the strategies. Throughout the GA exact arithmetics had to be used, which, since the GA has been implemented in Java, made use of the java.math.BigInteger class.

The GA itself has been set up to run for 250 rounds or until an optimal control had been reached, whichever would occur first. The size of the population has been set to 500 individuals. Both values could be increased to get a higher chance to reach an optimal solution. However, the computational cost of using exact arithmetics is quite high, which led to this acceptable compromise.

6.2.2. *Fitness Function.* In the example we used the conviction interval of $[0.375, 0.625]$, the most obvious fitness function to use is to count all the convinced individuals. This fitness function is denoted as *MaxVotes (MV)*. Figure 2 depicts a typical run of the GA with the MV fitness function.

While the MV fitness function provides an exact mapping of the original problem to the formulation of the GA's fitness function, it does not provide a very good example of a fitness function for a GA because it is discontinuous and has huge gaps in its range. Because the fitness function is just a count of the convinced individuals, the function can take integer values only. However, the GA's fitness function is defined on real values. As a consequence, the fitness

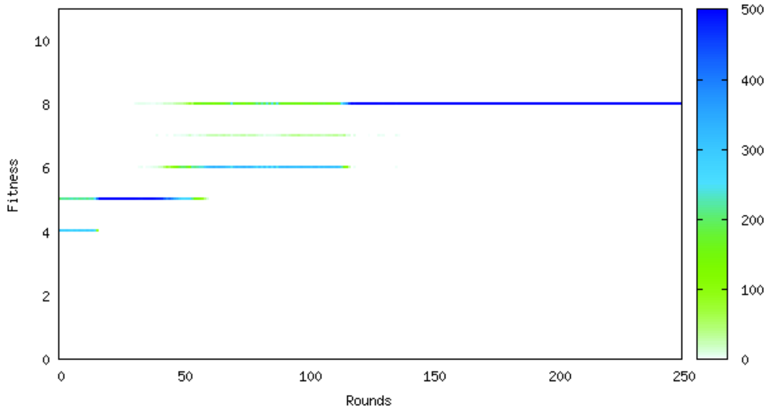

FIGURE 2. GA with the *MaxVotes* fitness function yielding at most eight voters

function does not provide enough information about a particular strategy. In the figure, one can see that, approximately after round 120, almost all remaining strategies yield eight voters, but the GA does not progress further to yield more voters. This indicates a lack of information on the strategies. If all chromosomes evaluate to the same fitness value, the GA has no way to select the fittest chromosomes and advance them to the next generation.

Therefore, we came up with some different fitness functions that span the entire range and try to provide additional information about strategies that yield the same amount of voters. These strategies, however, deviate from the original problem and create a slightly different problem for the GA to solve. Thus, all fitness functions that do not map the original problem directly, have to be evaluated with respect to their mapping ability.

The fitness functions fall into three different categories:

Weighted Sum: This category of fitness functions calculates the weighted sum of all the individuals final positions. The weight to be used is computed with a given partially defined function that maps a position to a weight. The MV fitness function is a special case of the Weighted Sum class of functions, because it assigns weight 1 to all positions in the conviction interval and weight 0 otherwise. The other functions used in this class are *DistanceToParty2 (D2P2)* and *BorderDistanceToAll (BD2A)*. Both of them differ from MV in that they assign values between 0 and 1 to positions that (a) are not in the conviction interval with higher values the closer the position is to the interval or (b) decreasing values within the interval, the closer the position is to the center of the interval. This leads to positions of individuals on the very edges of the interval to be the most favorable. The idea behind evaluating positions within the interval differently is that individuals sitting on the edges of the interval have the greatest effect on individuals that are not in the range yet.

Last Remaining: The class of the last remaining fitness functions does not evaluate every individual but restricts itself to convinced individuals (counted with weight 1) and the nearest individual that has not reached the conviction range yet (weighted according to the function). All other individuals are assigned weight 0. The fitness function in this category is *BorderDistanceToMin (BD2M)*, which has the same form as BD2A from the *Weighted Sum* category.

Minimum Distance: The last class of fitness function does not evaluate all individuals positions but takes into account the distance between the two outmost individuals (i. e.

the individual with the highest opinion and the individual with the lowest opinion). Because of the order preserving characteristics of the model, these two individuals do not change throughout one run, which means, we can just use the distance between the individual that started with opinion 0 and the one that started with opinion 1. If the distance is in the range of $[0, 0.25]$, which is the size of the conviction interval, the fitness evaluates to the maximum fitness value of 10. If the distance is greater than 0.25, the function computes a value that is decreasing to 0 with increasing distance. The two functions used in this class are *MinimumDistanceBetweenFirstAndLast (MDBFL)* and *MinimumDistanceBetweenFirstAndLastSquare (MDBFLS)* that have a linear or quadratic slope respectively. A third function in this class is the *MinimumDistanceBetweenFirstAndLastToCenterSquare (MDBFL2CS)* that accounts for the fact that the group of individuals with opinions below 0.5 may not behave symmetrically to the group with opinions above 0.5. This asymmetry can result in the position range not to be centered around 0.5 but deviate from that midpoint. Such behavior is undesirable, since the original goal is to get as many individuals close to 0.5 as possible. MDBFL2CS accounts for this and evaluates the positions of the two outmost individuals with respect to their distance to the desired midpoint. The two values are added.

6.2.3. *Evolution.* After the evaluation phase of the GA, the fittest chromosomes are chosen to advance to the next generation. There are different selection algorithms available. We used two different ones, which are among the standard selection algorithms:

Weighted Roulette Selector (WRS): Each chromosome is assigned a probability to advance to the next round proportional to its fitness. Then the population for the next round is chosen by randomly picking a chromosome from the so called "roulette wheel" as often as desired. This selection method allows for some chromosomes with low fitness values to advance to the next round, which results in a lower chance to reach a local optimum too quickly.

Best Chromosomes Selector (BCS): The BCS sorts the population according to the fitness values and discards the fraction with the lowest fitness values. The ration of chromosomes to retain is configurable. BCS fosters depth search with the danger of reaching a local optimum. As an advantage, it progresses much quicker than WRS.

After the chromosomes for the next round have been selected, the GA performs the crossover and mutation operations, whose parameters (percentage of the mutation, point of crossover) are configurable.

6.2.4. *Results.* Figure 3 show the performance of the different fitness functions. One notable observation is the step like behavior of the MV fitness function: Already around round 30, it reaches eight individuals, but does not advance from there. All the other fitness functions show a much smoother behavior. However, with the exception of BD2A and MDBFLS, all functions seem to have reached a plateau around round 150.

Judged from the performance of the fitness functions, BD2A promises the best results as it still progresses at round 250 and also reached a high fitness value. However, as BD2A optimizes a slightly different problem than the original problem. Therefore, the performance with respect to the fitness function has to be compared with the performance of the fitness function with respect to the voters.

Figure 4 depicts the performance of the fitness functions MV, BD2M, and MDBFLS. While MV maps the original problem exactly, only BD2M provides a good mapping. MDBFLS (and all other fitness functions alike) does not provide a good mapping of a fitness value to a certain number of voters. This, of course, poses a problem, since the only way for the GA of evaluating

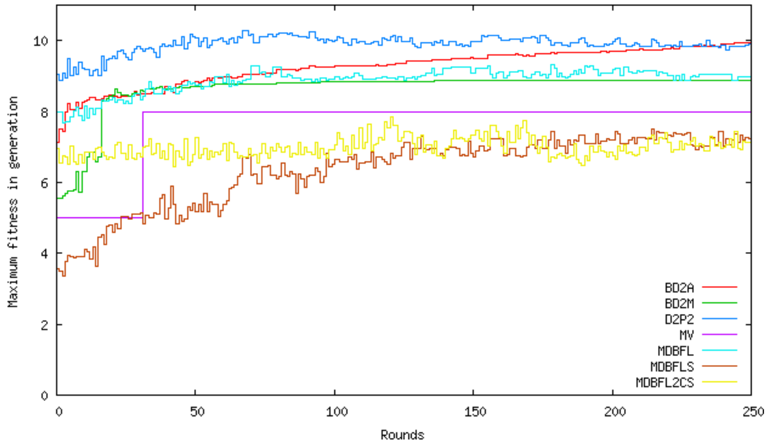

FIGURE 3. Different fitness functions reach different results

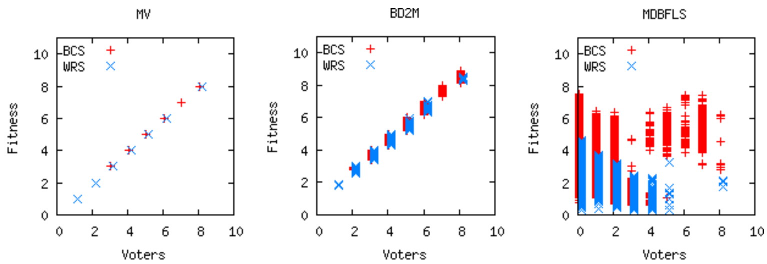

FIGURE 4. Performance of three fitness functions with respect to the convinced individuals

a certain strategy is the fitness function. The graph suggest that, apart from MV, only BD2M should be used.

Of the two different selectors available, BCS proofed to be the most useful of the two. While the WRS worked, the evolution of the population happened very slowly regardless of the fitness function used. Figure 5 shows the result of the two selectors while using the same fitness function (BD2A). Similar results hold for all other fitness functions as well.

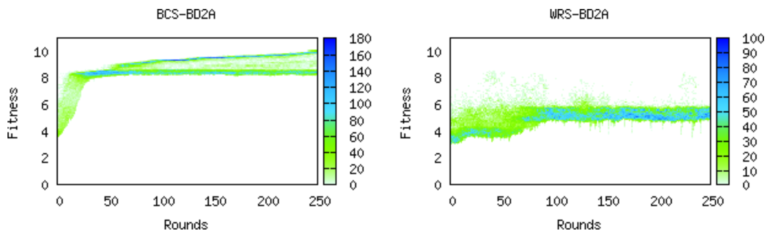

FIGURE 5. Comparison of two different selectors

Because of these findings, the BD2M fitness function has been tested with different values for the best performing selector BCS. The selector allows to configure the fraction of individuals that advances from one generation to the next. With a value of 50% only the better half of chromosomes advances. This leads to an extremely narrow search that runs into high risks of lingering at a local optimum. To increase the chances of leaving a local optimum again, the percentage should be increased. In the simulation, runs with ratios from the set $\{0.5, 0.6, 0.7, 0.75, 0.80, 0.85, 0.9, 0.91, 0.92, 0.93, 0.94, 0.95, 0.96, 0.97, 0.98, 0.99\}$ have been used. From these runs, only ratios above 0.75 resulted in stable evolution patterns, while rates below had their fitness values alternate between very low and very high values but did not converge.

Above 75% all runs converged with an optimal rate at around 95%. The runs with ratios of 0.95 and 0.96 were the ones that produced strategies that at least gave nine convinced individuals. Figure 6 shows these two runs and the number of voters that each chromosome generated.

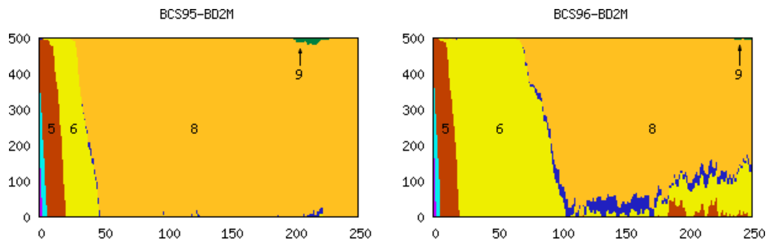

FIGURE 6. The two runs for different values for the BCS survival rate

Altogether the GA provides a heuristic to find optimal (or near optimal) solutions. However, because of the problem structure the GA did not find an optimal solution for the control problem. The most convinced individuals that the GA could find a strategy for were nine in the example setting. Strategies yielding nine individuals were extremely rare and could only be obtained in settings with highly tuned parameters. This result seems to indicate that the solution space has a very sparse population that could possibly occupy a very restricted region in the space.

6.3. **Receding horizon.** (Jörg)

7. Conclusion and outlook

Extensions:

- Investigation of chaotic behavior of the bounded-confidence dynamics
- Each party has its own controller (game theory)
- Multi-dimensional opinion space

7.1. **A game theoretic point of view.**

7.2. **Generalizations of the opinion space.**

References

1. Jan Christian Dittmer, *Consensus formation under bounded confidence.*, Nonlinear Anal., Theory Methods Appl. **47** (2001), no. 7, 4615–4621 (English).
2. S. Fortunato, *The Krause-Hegselmann Consensus Model with Discrete Opinions*, International Journal of Modern Physics C **15** (2004), 1021–1029.
3. ———, *Damage spreading and opinion dynamics on scale-free networks*, Physica A Statistical Mechanics and its Applications **348** (2005), 683–690.
4. ———, *On the Consensus Threshold for the Opinion Dynamics of Krause-Hegselmann*, International Journal of Modern Physics C **16** (2005), 259–270.
5. S. Fortunato and D. Stauffer, *Computer Simulations of Opinions*, ArXiv Condensed Matter e-prints (2005).
6. Santo Fortunato, *The Krause-Hegselmann consensus model with discrete opinions.*, Int. J. Mod. Phys. C **15** (2004), no. 7, 1021–1029 (English).

7. ———, *On the consensus threshold for the opinion dynamics of Krause-Hegselmann.*, Int. J. Mod. Phys. C **16** (2005), no. 2, 259–270 (English).
8. ———, *The Sznajd consensus model with continuous opinions.*, Int. J. Mod. Phys. C **16** (2005), no. 1, 17–24 (English).
9. Santo Fortunato, Vito Latora, Alessandro Pluchino, and Andrea Rapisarda, *Vector opinion dynamics in a bounded confidence consensus model.*, Int. J. Mod. Phys. C **16** (2005), no. 10, 1535–1551 (English).
10. Rainer Hegselmann and Ulrich Krause, *Opinion dynamics and bounded confidence: models, analysis and simulation*, Journal of Artificial Societies and Social Simulation **5** (2002), no. 3.
11. ———, *Truth and cognitive division of labour: First steps towards a computer aided social epistemology*, Journal of Artificial Societies and Social Simulation **9** (2006), no. 3.
12. Ulrich Krause, *Time-variant consensus formation in higher dimensions.*, Elaydi, Saber (ed.) et al., Proceedings of the 8th international conference on difference equations and applications (ICDEA 2003), Masaryk University, Brno, Czech Republic, July 28–August 1, 2003. Boca Raton, FL: Chapman & Hall/CRC. 185-191 (2005)., 2005.
13. ———, *Arithmetic-geometric discrete systems.*, J. Difference Equ. Appl. **12** (2006), no. 2, 229–231 (English).
14. Wim Liebrand, Andrzej Nowak, and Rainer Hegselmann (eds.), *Computer modeling of social processes*, Sage Publications, 1998.
15. J. Lorenz, *A stabilization theorem for dynamics of continuous opinions*, Physica A Statistical Mechanics and its Applications **355** (2005), 217–223.
16. Jan Lorenz, *Consensus strikes back in the hegselmann-krause model of continuous opinion dynamics under bounded confidence*, Journal of Artificial Societies and Social Simulation **9** (2006), no. 1.
17. K. Malarz, *Truth seekers in opinion dynamics models*, ArXiv Physics e-prints (2006).
18. F.A. Rodrigues and L. Da F.Costa, *Surviving opinions in Sznajd models on complex networks.*, Int. J. Mod. Phys. C **16** (2005), no. 11, 1785–1792 (English).
19. D. Stauffer, *The Sznajd Model of Consensus Building with Limited Persuasion*, International Journal of Modern Physics C **13** (2002), 315–317.
20. ———, *How to Convince Others? Monte Carlo Simulations of the Sznajd Model*, AIP Conf. Proc. 690: The Monte Carlo Method in the Physical Sciences (J. E. Gubernatis, ed.), November 2003, pp. 147–155.
21. ———, *Sociophysics simulations II: opinion dynamics*, AIP Conf. Proc. 779: Modeling Cooperative Behavior in the Social Sciences (P. Garrido, J. Maroo, and M. A. Muñoz, eds.), July 2005, pp. 56–68.
22. D. Stauffer and M. Sahimi, *Discrete simulation of the dynamics of spread of extreme opinions in a society*, Physica A Statistical Mechanics and its Applications **364** (2006), 537–543.
23. G. Weisbuch, *Bounded confidence and social networks*, The European Physical Journal B - Condensed Matter and Complex Systems **38** (2004), no. 2, 339–343.

Torsten Eymann, Fakultät für Rechts- und Wirtschaftswissenschaften, Universität Bayreuth, Germany
E-mail address: `torsten.eymann@uni-bayreuth.de`

Rainer Hegselmann, Fakultät für Kulturwissenschaften, Universität Bayreuth, Germany
E-mail address: `rainer.hegselmann@uni-bayreuth.de`

Stefan König, Fakultät für Rechts- und Wirtschaftswissenschaften, Universität Bayreuth, Germany
E-mail address: `stefan.koenig@uni-bayreuth.de`

Sascha Kurz, Fakultät für Mathematik, Physik und Informatik, Universität Bayreuth, Germany
E-mail address: `sascha.kurz@uni-bayreuth.de`

Christoph Niemann, Fakultät für Rechts- und Wirtschaftswissenschaften, Universität Bayreuth, Germany
E-mail address: `christoph.niemann@uni-bayreuth.de`

Jörg Rambau, Fakultät für Mathematik, Physik und Informatik, Universität Bayreuth, Germany
E-mail address: `joerg.rambau@uni-bayreuth.de`