

OPTIMAL OPINION CONTROL: THE CAMPAIGN PROBLEM

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

ABSTRACT. Opinion dynamics is nowadays a very common field of research. In this article we formulate and then study a novel, namely strategic perspective on such dynamics: There are the usual ‘normal’ agents that update their opinions, for instance according to the well-known bounded confidence mechanism. But, additionally, there is at least one strategic agent. That agent uses opinions as freely selectable strategies to get control on the dynamics: The strategic agent of our benchmark problem tries, during a campaign of a certain length, to influence the ongoing dynamics among normal agents with strategically placed opinions (one per period) in such a way, that, by the end of the campaign, as much as possible normals end up with opinions in a certain interval of the opinion space. Structurally, such a problem is an optimal control problem. That type of problem is ubiquitous. Resorting to advanced and partly non-standard methods for computing optimal controls, we solve some instances of the campaign problem. But even for a very small number of normal agents, just one strategic agent, and a ten-period campaign length, the problem turns out to be extremely difficult. Explicitly we discuss moral and political concerns that immediately arise, if someone starts to analyze the possibilities of an “optimal opinion control”.

1. INTRODUCTION

Since about 60 years the dynamics of opinions has been studied. Today it is a standard topic of general conferences on agent-based modelling. A bunch of models were defined and analyzed.¹ In the last 15 years at least hundreds and probably more than thousands of simulation studies on the dynamics of opinions were published.² The studies and their underlying models differ in many details: The opinions and the underlying time are continuous or discrete, the updating of opinions is governed by different updating regimes, the space of possible opinions may have more than one dimension, the dynamics may run on this or that type of network, various types of noise may be involved. But despite of all their differences there is a commonality in all these studies and their models: The agents influence their opinions mutually, *but they do not do that strategically*.

What is studied in the huge body of articles is typically focusing on convergence, dynamical patterns or final structures. Given the specific parameters of the underlying model, the typical questions are: Under what conditions does the dynamics stabilize? Does the dynamics lead to consensus, polarisation, or other interesting types of clustering? What are the time scales that are involved? What remains unasked in these studies are strategic questions like: Which

Key words and phrases. Opinion Dynamics, Optimal Opinion Control, Bounded Confidence, Mixed Integer Linear Programming, Heuristics.

¹For the general history of opinion dynamics see the introduction and for a partial classification see ch. 2 and 3 of the paper by Hegselmann & Krause (2002).

²There are several surveys encompassing various models of opinion dynamics (Acemoglu & Ozdaglar 2011; Castellano *et al.* 2009; Liebrand *et al.* 1998; Stauffer 2005; Xia *et al.* 2011). A microfoundation for the evolution of several opinion dynamics mechanisms is proposed and discussed by Groeber *et al.* (2014). – In subsection 1.3 we give specific hints to the related opinion dynamics literature that is directly relevant in the context of our present article.

opinion should an agent pretend to have to drive the whole dynamics in his or her preferred direction? Where in the opinion space should an agent ‘place’ an opinion, given that he or she has a certain preference with regard to the opinions in later periods. Our article deals with such strategic questions. We develop a conceptual framework that allows to answer strategic questions in certain cases. Additionally, we analyze why it is surprisingly difficult or impossible to give exact answers to strategic questions even in cases that look very, very simple.

It is not by accident that strategic questions are normally not raised in the sort of research that is labeled opinion dynamics. The standard approach is to describe the dynamics as a more or less complicated dynamical system: There is a set I of agents $1, 2, \dots, i, j, \dots, n$ and a discrete time $t = 1, 2, \dots$. The opinions of the agents are given by an *opinion profile*. This is a vector $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$ that describes the state of the system at time t . Even if stated in an informal or semi-formal way (sufficiently clear to program the process), the dynamics of the system is basically given by a function \mathbf{f}^t that computes the state of the system \mathbf{x}^{t+1} as $\mathbf{x}^{t+1} = \mathbf{f}^t(\mathbf{x}^t)$.

Thus, for each agent i the function \mathbf{f}^t specifies how x_i^{t+1} depends upon \mathbf{x}^t . Depending upon the specific opinion dynamics model, the vector valued functions \mathbf{f}^t work in very different ways. For the most part they do some kind of averaging: averaging with a privileged weight for agent i 's own opinion or a weighted averaging with weights w_{ij} that agent i assigns to agents j and that are meant to express agent i 's respect for agent j , or some other sort of averaging subject to constraints, for instance constraints in terms of network distance on an underlying network on which the opinions dynamics is assumed to run.

Whatever the ‘story’ about \mathbf{f}^t , it is always a reactive process in which the agents react on the *last* period t . In principle the step to $t + 1$ might depend upon some more past periods. But even then, an answer to the question, where to place an opinion in order to drive the dynamics in a preferred direction, requires something very different from looking into the past: It requires *anticipation*, i.e., finding out what the future effects of placing an opinion here or there in the opinion space probably are, and then placing it there, where the placement is most effective to get to a preferred outcome.

In the following, we assume a setting in which we have two sets of agents: First, a set of *non-strategic* agents as they are usually assumed in opinion dynamics. They are driven by the function \mathbf{f}^t . The function describes a dynamical system in which the non-strategic agents always reveal their ‘true’ actual opinion, take the opinions of others as their true actual opinion, and mutually react on the given opinion profile \mathbf{x}^t according to $\mathbf{f}^t(\mathbf{x}^t)$. The second set of agents is a set of *strategic* agents. Whatever their true opinion may actually be, they can place any opinion strategically in the opinion space where, then, non-strategic agents take these opinions as their face values and consider them as revealed true opinions of other agents. The strategic agents have preferences over possible opinion profiles of non-strategic agents. Therefore, they try to place opinions in such a way that the opinion stream that is generated by \mathbf{f}^t is driven towards the preferred type of profile.

Our setting has a structure as it is often perceived or even explicitly ‘conceptualized’ by political or commercial campaigners: There is an ongoing opinion stream, result of and driven by mutual exchange of opinions between communicating citizens, consumers, members of some parliamentary body, etc. That opinion stream has its own dynamics. However, it is possible to intervene: Using channels of all sorts (TV, radio, print media, mails, posters, ads, calls, speeches, personal conversation, etc.) one can place opinions in the opinion space. Done in an intelligent way, these interventions should drive the opinion stream in the direction of an outcome that is preferred by the principal who pays a campaign. About that will often be the self-understanding and selling-point of a campaigning agency.

The number of strategic agents matters: If there are *two or more* strategic agents, the setting becomes a *game theoretical context* in which the strategic agents have to take into account that

there are others that try to influence the opinion dynamics *as well*. Therefore, the strategic agents do not only ‘play a game’ with the non-strategic agents. They play also – and that now in an exact game theoretical sense of the word – a *game* against each other. It is a complicated game for which in principle usual solution concepts like the (sub-game perfect) Nash equilibrium can be applied. But if there is *just one* strategic agent, then there are no other competing players. That turns the problem of the strategic agent into the following question: How can one place opinions in an ongoing opinion stream (governed by \mathbf{f}^t) in such a way that the stream moves as much as possible in the direction of one’s favorite profile? This task means to optimize of decisions over time. Technically speaking, problems of this type are *optimal control problems*. As soon as a mathematical formalization is available, there are various mathematical methods to find solutions.

The topic of this paper is the optimal control problem of a strategic agent who tries to influence a certain ongoing opinion dynamics. As a benchmark problem for such an agent, we define what we call the *campaign problem*: The strategic agent tries to control an opinion dynamics, such that in a certain period, known in advance, there are as much as possible opinions of normal agents in a certain interval of the opinion space. That covers many types of voting or buying campaigns. We will investigate the campaign problem by various methods. Our focus will always be on understanding basic features of the problem—and that will be difficult enough. Some elementary mathematical proofs will give some theoretical insights into structural properties. With models from mixed-integer linear programming (MILP) we will try to directly solve our central optimization problem. This way we can solve some instances of the campaign problem, but – surprisingly – not all. Therefore we will additionally attack the control problem by heuristic methods. An additional investigation using genetic algorithms provides evidence for the fact that ‘good’ controls are almost impossible to find by randomized exploration.

1.1. Moral and political concerns. Our question and approach immediately raises moral and political concerns: Over the last years we learnt that there are well equipped agencies that (among other things) aim at a more or less complete supervision of private and public opinions. Isn’t the title of our paper at least a partial confession, that now, as a kind of additional thread, basic research on strategies for an efficient manipulation of public opinion formation is put on the agenda, isn’t our ‘control terminology’ a tell-tale language?

This is a very serious question and a very serious concern. An answer requires careful consideration of least six aspects: *First*, the only means of opinion ‘control’ in our setting is publically stating an opinion, which should not be confused with opinion manipulation based on arbitrary methods. *Second*, campaigning can be done for both good and bad purposes: Information, enlightenment, spreading the truth, *or* disinformation, confusion, spreading lies. *Third*, whoever plans a campaign for a good purpose will immediately get into serious optimization problems. For an example, let’s assume that measles vaccination is basically a good thing. From a public health point of view, a certain minimum vaccination rate is necessary and a certain upper rate sufficient. In the US, Germany, and other countries the actual vaccination rates are too low. Therefore public health institutions design pro vaccination campaigns. In doing so, important questions are: On which parts of the networks of vaccination skeptics and enemies should one concentrate to what degree, or in what sequence? What might convince which parts of the networks to change their opinions? Which opinion changes can be induced in the network? – All these questions are questions about an optimal campaign, given one’s constraints (in terms of budget, time, channels, arguments, their effects on whom, and chances to confront people with them etc.). *Fourth*, our control terminology is the usual terminology in the disciplines, theories, and tool box approaches that can be used to solve optimization problems of all sorts. Our campaign problem is just an instance. The spirit of these approaches is technical: One has some

constrained means to influence something to some degree in a certain, more or less attractive direction. What, then, is an optimal use of one’s means?—that is the guiding question. The question is neutral as to ends, except for optimality itself (and some more formal consistency requirements). The possibility of a ‘dual use’ is inherent to such approaches. *Fifth*, in what follows, our strategic agents place their opinions independent of any truth considerations. They try to influence the ongoing opinion dynamics in their favorite direction by placing an opinion here or there in the opinion space. Only the effects matter. Therefore, one might say, our strategic agents are completely opportunistic. Under a less pejorative description we might consider them as perfect diplomats that know how to overcome entrenched opinions. If we consider diplomacy or opportunism as morally unacceptable, then we could formulate a corresponding *moral* constraint: We could require to place only opinions that are close to one’s true actual opinion, and we specify ‘close’ by a certain threshold for the maximal acceptable distance of placed opinions to one’s own actual opinion. Probably nobody would advocate a threshold of zero.³ But, then, one should also recognize, that for any such *non-zero*-threshold, there still exists the control problem of how to design an optimal campaign (though with an additional constraint). *Finally*, dual use concerns about an optimal control approach to opinion dynamics are justified, of course. The existence of professional disinformation agencies is not an invention of conspiracy theorists. They are at work, they have their *secret* expertise – and, in all likelihood, they learn. But exactly because of that there has to be *public* knowledge about what can be done in terms of optimal campaigning – and that for both, good *and* bad purposes. The first type of knowledge is supportive, the second protective. Helpful are both.

1.2. Our benchmark: The campaign problem. To specify the campaign problem we add *just one* strategic agent₀ to the set of agents. Agent₀ is equipped with the ability to freely choose at any time step what other agents then perceive as his or her opinion. We call agent₀ the *controller* and his freely chosen opinion the *control*. Mathematically, this makes the development of the opinion system dependent on exogenously chosen parameters, namely the control opinion, and we are faced with a *control system*. If we define what the controller wants to achieve, we can formulate an *optimal control problem*, for which the controller tries to find the controls to get there.

Our optimal control problem is of a *specific – seemingly simple – type*: Agent₀ can strategically place opinions in a finite sequence of periods, one and only one opinion per period. There is the ongoing and underlying opinion dynamics, given by a function \mathbf{f}^t . Period by period agent₀ tries to do the placement of an opinion in such a way that finally in a certain future time step N (the *horizon*), known in advance, something is maximized: the number of normal agents’ opinions that are in a certain part of the opinion space that *ex ante* was specified by agent₀. To keep it simple, we assume as a special one-dimensional opinion space the real-valued unit interval $[0, 1]$. As the part of the opinion space preferred by the controller, we assume a *target interval* $[\ell, r] \subseteq [0, 1]$ for some prescribed $\ell < r$ in $[0, 1]$ known to the controller.

Both assumptions are much less restrictive than they seem to be: First, the unit interval can be used to represent opinions about, for example, tax rates, minimal wages, maximum salaries, political positions on a left-right spectrum, product quality, or any property whatsoever that can be expressed by real-valued numbers. If – and often that will be the case – the ‘really’ possible range of numerical values is different from the unit interval, then some transformation, approximation, or range regarding guess work is necessary. But that is an easy and widely accepted step (at least in principle). Second, suppose there are m fixed alternatives $\mathbf{a} = (a_1, a_2, \dots, a_m) \in [0, 1]^m$,

³ Imagine, I use a hypothetical argument that I myself do not share. Nevertheless, in other persons’ opinions this may induce a change in the direction of what I consider the truth. What is morally wrong with such a discussion strategy?

sorted such that $a_1 \leq a_2 \leq \dots \leq a_m$. Further suppose, our n normal agents have to choose among the alternatives at the future time step N and will do that by choosing an alternative that is next to their own opinion in that time step. What, then, is the problem of a controller with the interest to make as much normal agents as possible choosing a certain alternative a_j ? Obviously the problem is to maximize the number of agents' opinions that in time step N are within a certain interval to the left and to the right of the favored alternative a_j . The exact bounds of that interval depend upon the exact positions of the two nearest alternatives a_{j-1} to the left and a_{j+1} to the right of the favored a_j . The exact left and right bounds are then $\frac{a_{j-1}+a_j}{2}$ and $\frac{a_j+a_{j+1}}{2}$, respectively.

Therefore, whatever the vector of alternatives may be (e.g., positions of parties or candidates on a left/right scale, a price or a technical specification of a certain type of product), whenever there are voting or buying decisions⁴ after a foregoing opinion dynamics (e.g., about the appropriate position in the political spectrum, the acceptable price or a desirable technical specification of some product), our controller agent₀ who tries to 'sell' a certain alternative a_j as effectively as possible, has always the same problem: How to get by period N as many opinions as possible within a certain target interval $[\ell, r]$, determined by the closest left and right competitors of a_j – Obviously, our framework and setup is much more general than it looks at a first glance.

1.3. The underlying opinion dynamics: A linear and a non-linear version. Our problem and approach presupposes an underlying opinion dynamics given by a function \mathbf{f}^t . But there are many. We will use a linear and a non-linear variant, each of them probably being the most prominent variant of their type.

In the *linear variant*, the opinion dynamics is driven by *weighted averaging*: Weights w_{ij} may express the respect or competence that an agent i assigns to an agent j ; alternatively, a weight w_{ij} may express the influence, impact or power that agent j has on agent i . The weights that an agent assigns sum up to one. The opinion space is the unit interval $[0, 1]$. – The history of this linear system goes back to French Jr (1956), it has been already presented by Harary (1959), it was explicitly stated by DeGroot (1974), and it received a lot of attention, especially in philosophy, through the book by Lehrer & Wagner (1981).⁵ We will refer to that model as the *DeGroot-model* (DG-model).⁶

The *non-linear variant* that we will use is the so called *bounded confidence model* (BC-model). In this model the agents take seriously those others whose opinion are not too far away from their own opinion: The agents have a certain confidence radius ϵ and update their opinions – the opinion space is again the unit interval $[0, 1]$ – by averaging over all opinions that are from their own opinion not further away than ϵ : An agent i updates to x_i^{t+1} by averaging over the elements

⁴It may even be a buying decision in the metaphorical sense of whether or not to 'buy' a certain assumption: Imagine a committee that after some discussion has to decide whether to proceed based on this or that assumption in a vector \mathbf{a} of alternatives.

⁵Lehrer & Wagner (1981) do not interpret the iterated weighted averaging as a process in time. As stressed by Hegselmann & Krause (2006, p. 4): "Their starting point is a 'dialectical equilibrium', i.e., a situation after, the group has engaged in extended discussion of the issue so that all empirical data and theoretical ratiocination has been communicated. '... the discussion has sufficiently exhausted the scientific information available so that further discussion would not change the opinion of any member of the group' ((Lehrer & Wagner 1981, p. 19)). The central question for Lehrer and Wagner then is: Once the dialectical equilibrium is reached, is there a rational procedure to aggregate the normally still divergent opinions in the group (cf. (Lehrer 1981, p. 229)? Their answer is 'Yes.' The basic idea for the procedure is to make use of the fact that normally we all do not only have opinions but also information on expertise or reliability of others. That information can be used to assign weights to other individuals. The whole aggregation procedure is then iterated weighted averaging with $t \rightarrow \infty$ and based on constant weights. It is shown that for lots of weight matrices the individuals reach a consensus whatever the initial opinions might be – if they only were willing to apply the proposed aggregation procedure."

⁶In philosophy the model is often called the *Lehrer-Wagner model*.

of the set $\{j \in \{1, 2, \dots, n\} \mid |x_i^t - x_j^t| \leq \epsilon\}$, i.e., over all the opinions that are within what is called his or her confidence interval. The model was defined by Krause (1997), in 1998 it was coined *bounded confidence model* (Krause 2000), and then for the first time, to a certain extent, comprehensively analyzed, by both simulations and rigorous analytical means, by Hegselmann & Krause (2002).

The model looks extremely simple. However, there are several warnings in the literature on the BC-model, among them a very recent one: “The update rule . . . is certainly simple to formulate, though the simplicity is deceptive” (Wedin & Hegarty 2014, p. 2). The authors’ warning is well founded: The simple update rule generates a complicated dynamics that still is only partially understood. The main reason for that is this: The updating rule of the BC-model can be described as assigning weights to other agents. All agents with opinions out of the confidence interval get a weight of 0; agents within get a weight of 1 divided by the number of agents that are within the interval. Therefore, the BC-dynamics is weighted averaging as well. However, there is a crucial difference to the linear DG-model: The weights of the BC-model are *time-dependent* and, even worse, *discontinuously dependent on the current profile*. That causes a lot of trouble – and, at the same time, generates many of interesting effects. As a consequence, the BC-model became a subject for all sorts of analytical or computational analysis and a starting point for extensions of all sorts. The body of literature on the BC-model is correspondingly huge.

Structural results in the BC-model were obtained with respect to convergence and its rate (Dittmer 2001; Krause 2006; Lorenz 2006), thresholds for the confidence radius (Fortunato 2005b), the identification of the really crucial topological and metric structures (Krause 2006), or the influence of the underlying network (Weisbuch 2004). The influence of a ‘true’ opinion, to which (some of) the individuals are attracted, received special attention (Hegselmann & Krause 2006; Malarz 2006; Douven & Riegler 2010; Douven & Kelp 2011; Kurz & Rambau 2011; Wenmackers *et al.* 2012, 2014). With a grain of salt, the true opinion can also be interpreted as a control that is constant over time and that is contained in each individual’s confidence interval.

Many variants of the original BC-model (discrete time, finitely many individuals, continuous opinion space) have been proposed, among them pairwise sequential updating (Deffuant *et al.* 2000), a discrete opinion space (Fortunato 2004), a multi-dimensional opinion space (Fortunato *et al.* 2005; Krause 2005), a noisy opinion space (Bhattacharyya *et al.* 2013; Pineda *et al.* 2013), a continuous distribution instead of a finite-dimensional vector of opinions (Lorenz 2005, 2007), continuous time (Blondel *et al.* 2010), and a continuum of agents (Wedin & Hegarty 2014). Alternative dynamics have been enhanced with BC-type features by several authors (Stauffer 2002, 2003; Fortunato 2005; Stauffer & Sahimi 2006; Rodrigues & Da F.Costa 2005) in order to make the resulting emergent effects more interesting.

It is interesting to note that simulations play an important role in the investigation of the innocent-looking models for opinion dynamics (Stauffer 2003; Hegselmann & Krause 2006; Fortunato & Stauffer 2006; Stauffer & Sahimi 2006) – a hint that some aspects of opinion dynamics are very hard to deduce purely theoretically. The arguably most general account of theoretical aspects was contributed by Chazelle on the basis of a completely new methodology around function theoretical arguments (Chazelle 2011). Moreover, opinion dynamics in multi-agent systems can be seen as an instance of *influence system* – this broader context is also described by Chazelle (2012).

Carletti *et al.* (2006) study the influence of a given, fixed exogenous propaganda opinion on the Deffuant-Weisbuch opinion dynamics. However, the problem of optimally choosing a propaganda opinion is not considered. Implicitly, the control of opinion dynamics started out from gaining control over the communication structure (Lorenz & Urbig 2007) or additional system parameters (Lorenz 2008). Recently, the effect of time-varying exogenous influences on the consensus process has been studied by Mirtabatabaei *et al.* (2014). Kurz (to appear) investigated how the time

to stabilization changes if controlled opinions enter the scene. For the first time to the best of our knowledge, the notion of opinion control appeared literally only in the title of a paper by Zuev & Fedyanin (2012), based on a different dynamical model, though. For continuous time optimal control techniques have been applied to opinion consensus, exemplarily also for the BC-dynamics, in the preprint by Albi *et al.* (2014). Closest to our research is, up to now, probably the investigation by Fortunato (2005a) about so-called damage spreading: what happens if some of the opinions in a BC-model change erratically and possibly drastically? The setting in that paper, however, has not been utilized to find in some sense ‘optimal damages’, i.e., some that lead to an outcome that is most desirable among all possible outcomes.

1.4. Specification of the campaign problem for the DG- and the BC-model. In what follows, we analyze our optimal control problem with the linear DG- and the non-linear BC-model as the underlying opinion dynamics that agent₀ tries to control: Given the DG- or BC-dynamics and using the controls, i.e., placing one opinion per period here or there in the opinion space, agent₀ tries to maximize the number of agents with opinions that in a future period N are in a certain target interval $[\ell, r]$ – and therefore would ‘buy’ the corresponding alternative a_j that agent₀ is actually campaigning for. – In the following, we will switch to a more vivid language fitting this interpretation: agents are called *voters*, the special agent₀ is called the *controller*, the target interval $[\ell, r]$ is called the *conviction interval*, and voters in the target interval are called *convinced voters*.

As an example, we will analyze a specific instance of the campaign problem: There are 11 voters, governed by a DG- or BC-version of the (reactive) function \mathbf{f}^t . At $t = 0$, the 11 voters start with the opinions $0, 0.1, 0.2, \dots, 0.9, 1$. The confidence radius is given by $\epsilon = 0.15$ and the conviction interval is $[0.375, 0.625]$. That conviction interval would be the target in a campaign in which the alternatives $\frac{1}{4}, \frac{1}{2}, \frac{3}{4}$ are competing and $\frac{1}{2}$ is the preferred alternative. The goal is to maximize the number of convinced voters, i.e., those with opinions in the conviction interval, in a certain future period N with $N = 1, 2, \dots, 10$. The benchmark problem looks like a baby problem. But it is a monster: It will turn out that for higher one digit values of N we could not solve it by the most sophisticated methods available.⁷

More specifically, we will see, that even for this innocent-looking example we were not able to find the optimal number of convinced voters for all numbers of stages between 1 and 10. In Table 1 we summarize our results on the benchmark campaign problem in this paper. In the rest of the paper we will explain in detail how we obtained that knowledge.

1.5. Structure of this Paper. In what follows, we

- define exactly the new problem *optimal opinion control*⁸ and an example instance of it, the benchmark campaign problem, for illustration and test purposes,
- develop two exact mathematical models based on mixed integer linear programming to characterize optimal controls⁹,
- devise three classes of primal heuristics (combinatorial tailor-made, meta-heuristics, model predictive control) to find good controls,
- present computational results on all methods applied to our benchmark problem,

⁷These values were chosen in a research seminar because they looked very simple and came to our mind first. We were convinced that we would solve all related problems very fast by whatever method and could start to worry about more serious problem sizes. We were wrong. It is still this innocent-looking benchmark campaign problem that determines the agenda, as we will see.

⁸During finalizing work on this paper, it came to our attention that optimal opinion control was independently introduced and investigated with a different objective in continuous time by Wongkaew *et al.* (2015)

⁹Another application of mixed integer linear programming techniques in modeling social sciences was, e.g., presented by Kurz (2012).

# stages	# convincable voters	
	lower bound*	upper bound**
0	3	3
1	3	3
2	4	4
3	5	5
4	5	5
5	6	6
6	6	6
7	8	11
8	8	11
9	8	11
10	11	11

TABLE 1. The optimal number of convinced voters in the benchmark example lies between the given lower and upper bounds, depending on the number of stages; we will explain in in Sections 6* and 5** how we obtained this information.

- sketch possible lines of future research.

In Section 2 we formally introduce the optimal opinion control problem. Section 3 shows what can happen if the pitfalls of numerical mathematics are ignored in computer simulations. Section 4 presents some structural knowledge about optimal controls in small special cases, thereby indicating that a general theoretical solution of the optimal control problem is not very likely. Computational results are presented in Section 5. The detailed presentation of the underlying exact mathematical models is postponed to Appendix A, followed by information on the parameter settings for the commercial solver we used in Appendix B. Our heuristics are introduced in Section 6. In Section 7 we interpret the obtained results. Conclusions and further directions can be found in Section 8.

2. MODELING THE DYNAMICS OF OPINIONS AND THEIR CONTROL

We will now formalize the ideas sketched so far. In this paper, we will restrict ourselves to the arguably simplest case where an opinion of a voter $i \in I$ can be represented by a real number x_i in the unit interval $[0, 1]$.

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 voter $i \in I$ 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 $\mathbf{x}^{t+1} := \mathbf{f}^t(\mathbf{x}^t)$. We assume a given *start value* x_i^{start} for the current opinion of each voter $i \in I$. Thus, $x_i^0 = x_i^{\text{start}}$ holds for all $i \in I$.

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 DeGroot Model. In this model, each voter is again in contact with each other at every stage. The strengths of the influences of opinions on other opinions are given by non-negative weights w_{ij} with $\sum_{j \in I} w_{ij} = 1$ for all $i \in I$, with the meaning that the opinion of voter i is influenced by the opinion of voter 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:

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

Under mild conditions, it can be shown that this, in the limit, leads to consensus.¹⁰ It leads, as we will see below, still to an interesting optimal control problem.

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

$$I_i(x_1, \dots, x_n) := \{j \in I : |x_j - x_i| \leq \epsilon\}. \quad (2)$$

Observe that $I_i(x_1, \dots, x_n) \neq \emptyset$ due to $i \in I_i(x_1, \dots, x_n)$.

Then the system dynamics of the BC-model is given as follows:

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

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.

2.3. 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*.

Formally, this means in the simplest case (and we will restrict to this case) that the controller can present an additional opinion u^t at 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):

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

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

We can interpret this as a usual model of opinion dynamics with an additional opinion x_0 that can be positioned freely at every stage by the controller. The aim of the controller is 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 voters $j \in I$ with $x_j \in [\ell, r]$. We want to maximize the number

¹⁰One proof uses the Banach Fixed Point Theorem.

of convinced voters. Thus, the problem we want to investigate is the following deterministic discrete-time optimal control problem:

$$\begin{aligned} & \max_{x_0^0, x_0^1, \dots, x_0^{N-1}} |J(x_1^N, \dots, x_n^N)| \\ & \text{subject to} \\ & \quad x_i^0 = x_i^{\text{start}} \quad \forall i \in I, \quad (\text{Start Configuration}) \\ & \quad x_i^{t+1} = f_i(x_0^t; x_1^t, \dots, x_n^t) \quad \forall i \in I, t = 0, 1, \dots, N-1, \quad (\text{System Dynamics}) \\ & \quad x_0^t \in [0, 1] \quad \forall t = 0, 1, \dots, N-1, \quad (\text{Control Restrictions}) \end{aligned}$$

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

3. SIMULATION AND PITFALLS FROM NUMERICAL MATHEMATICS

No matter what we do: a computer can represent only finitely many distinct numbers. Thus, it is impossible that a computer can distinguish infinitely many numbers like there exist in $[0, 1]$. Even worse: if a number type is used for computer programming that uses a fixed number of bits, i.e., zeros and ones in the binary representation, like “float” or “double”, then distinct real numbers that are very close can be interpreted as identical numbers by the computer. This has a serious influence on our ability to check correctly whether or not one opinion is in the confidence interval of another whenever this is a cutting-edge decision. More specifically: If we try to solve our benchmark problem for the BC-model, the computer has – again and again – to decide the question whether or not $|x_i - x_j| \leq \epsilon$, where x_i , x_j , and ϵ are, in particular, *real* numbers. For a human being with a tiny bit of math training it is easy to answer the question whether $|0.6 - 0.4| \leq 0.2$. If a computer has to answer that simple question and uses what in some programming languages is called the data format “real” or “float” or even “double” (float with double precision), then the computer might get it wrong. In Figure 1, left, one can see that effect: We start the dynamics with 6 voters that are regularly distributed at the positions $0, 0.2, 0.4, \dots, 1$. The confidence radius ϵ is 0.2 .¹¹ Obviously the computer (using a program written in Delphi, but in NetLogo the analogous mistake would happen, possibly somewhere else) answers the question whether $|0.6 - 0.4| \leq 0.2$ in a wrong way. As a consequence, from the first update onwards the dynamics is corrupted: Given our start distribution and the homogeneous, constant, and symmetric confidence radius, the opinion stream should be mirror symmetric with respect to a horizontal line at $y = 0.5$. That symmetry is forever destroyed by the very first update. What happens here is no accident. It is the necessary consequence of the floating point arithmetic that computers use to approximate real numbers. Using floating point arithmetic each number is represented with a finite number of binary digits, so a small error is possibly made. For a hard decision like $|x_i - x_j| \leq \epsilon$ or $|x_i - x_j| > \epsilon$ a small error is sufficient to draw the wrong conclusion, whenever $|x_i - x_j|$ equals or is rather close to ϵ . The only way for us to cope with this problem is to resort to exact rational arithmetics throughout, although there may be more sophisticated methods to improve efficiency. This numerical instability has the more

¹¹The number 0.2 is tricky because its binary representation as a floating point number is determined by $0.2 = \frac{1}{8} + \frac{1}{16} + \frac{1}{128} + \frac{1}{256} + \dots = 0.\overline{0011}2$, i.e., its exact representation would need infinitely many bits. For example, a cutoff bit representation would be identical to a smaller number than 0.2. Thus, checking whether or not some number is smaller or larger than 0.2 checks this question not for 0.2 but for a smaller number. How exactly 0.2 is finitely represented in a computer is usually defined by an IEEE standard.

serious consequence that of-the-shelf optimization algorithms with floating point arithmetics can not be used without checking the results for correctness in exact arithmetics.¹²

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

$$\begin{aligned}
 x^0 &= (0.0, 0.2, 0.4, 0.6, 0.8, 1.0), \\
 x^1 &= (0.1, 0.2, 0.4, 0.6, 0.8, 0.9), \\
 x^2 &= (0.15, 0.2\bar{3}, 0.4, 0.6, 0.7\bar{6}, 0.85), \\
 x^3 &= (0.191\bar{6}, 0.26\bar{1}, 0.4\bar{1}, 0.5\bar{8}, 0.73\bar{8}, 0.808\bar{3}) \\
 &= \left(\frac{23}{120}, \frac{47}{180}, \frac{37}{90}, \frac{53}{90}, \frac{133}{180}, \frac{97}{120} \right), \\
 x^4 &= \left(\frac{163}{720}, \frac{311}{1080}, \frac{227}{540}, \frac{313}{540}, \frac{769}{1080}, \frac{557}{720} \right), \\
 x^5 &= \left(\frac{673}{2160}, \frac{673}{2160}, \frac{3271}{8640}, \frac{5369}{8640}, \frac{1487}{2160}, \frac{1487}{2160} \right), \\
 x^6 &= \left(\frac{577}{1728}, \frac{577}{1728}, \frac{577}{1728}, \frac{1151}{1728}, \frac{1151}{1728}, \frac{1151}{1728} \right) \\
 &= \left(0.3339120\bar{37}, 0.3339120\bar{37}, 0.3339120\bar{37}, 0.6660879\bar{62}, \right. \\
 &\quad \left. 0.6660879\bar{62}, 0.6660879\bar{62} \right).
 \end{aligned}$$

The corresponding correct trajectory is drawn on the right hand side of Figure 1.




FIGURE 1. A computational disaster caused by a tiny inaccuracy (left: numerical artefact; right: correct result)

As we have seen, a small error in the positions of the voters, computational or by observation, can have a drastic effect. We mention that this effect can not happen for numerical stable dynamics like, e.g., the DG-model. 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 other programming languages.

¹²Polhill *et al.* (2005) demonstrate that, based upon floating-point numbers, in several agent-based models branching statements lead to severe numerical artefacts.

There are quite a lot of articles dealing with the simulation of the BC-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 they are recalculated with exact arithmetic. For no existing publication, however, we found any evidence that the conclusions are only artefacts of numerical trouble. For results using randomized starting configurations the probability is zero that the distance between agents equals the confidence radius. In those experiments numerical artifacts are much less likely (though not impossible) than in simulations starting from equidistant configurations.

We have to admit, that in the starting phase of our investigation in optimal control of opinion dynamics, we have also used floating point arithmetic. We heuristically found controls achieving 10 voters after 10 stages. Using exact arithmetic it turned out that the computed control yields only 4 convinced voters, which is a really bad control, as we will see later on.

4. BASIC STRUCTURAL PROPERTIES OF OPTIMAL CONTROLS

In this section we collect some basic facts about structural properties of optimal controls, mainly for the BC-model. While generally the DG-model has nicer theoretical properties, there is an exception when considering the ordering of the voters over time.

Lemma 4.1. *(Cf. (Krause 2000, Lemma 2).) Consider an instance of the BC-model (with control or without).*

- (1) *If $x_i^t = x_j^t$, then $x_i^{t+1} = x_j^{t+1}$.*
- (2) *If $x_i^t \leq x_j^t$, then $x_i^{t+1} \leq x_j^{t+1}$.*

Proof. If the positions of voter i and voter j coincide at stage t , then they have identical confidence sets and the system dynamics yields the same positions for i and j at stage $t + 1$. For (2), we assume $x_i^t < x_j^t$ w.l.o.g. We set $C = I_i(x_1^t, \dots, x_n^t) \cap I_j(x_1^t, \dots, x_n^t)$, $L = I_i(x_1^t, \dots, x_n^t) \setminus C$, and $R = I_j(x_1^t, \dots, x_n^t) \setminus C$. Due to $x_i^t < x_j^t$ we have $x_i^t < x_c^t < x_r^t$ for all $l \in L$, $c \in C$, and $r \in R$. Thus $x_i^{t+1} \leq x_j^{t+1}$. \square

The analogous statement for the DG-model is wrong in general.

Next we observe that one or usually a whole set of optimal controls exists. The number of convinced voters is at any stage trivially bounded from above by the total number of voters $|I|$. Hence, to every control there corresponds a bounded integer valued number of convinced voters. With some technical effort an explicit bound can be computed.

First, we observe that there are some boundary effects. Consider a single voter with start value $x_1^0 = \frac{1}{2}$ and $\epsilon = \frac{1}{2}$. Further suppose that the conviction interval is given by $[0, \delta]$, where δ is small. The most effective way to move the opinion of voter 1 towards 0 is to place a control at 0 at each stage. With this we obtain $x_1^t = \frac{1}{2^{t+1}}$ for all t . Thus the time when voter 1 can be convinced depends on the length δ of the conviction interval. This is due to the fact that we can not place the control at $x_1^t - \epsilon$ if x_1^t is too close to the boundary. The same reasoning applies for the other boundary at 1. In order to ease the exposition and the technical notation we assume that no opinion is near to the boundary in the following lemma.

Lemma 4.2. *Consider an instance of the BC-model such that the start values and the conviction interval $[l, r]$ are contained in $[\epsilon, 1 - \epsilon]$. It is possible to select suitable controls at each stage such that after at most $\frac{2n+1}{\epsilon} + 2$ stages all $|I|$ voters are convinced.*

Proof. We will proceed in two steps. In the first step we ensure that all voters have the same opinion after a certain amount of stages. In the second step we will move the coinciding opinions inside the conviction interval.

Without loss of generality, we assume the ordering $x_1^0 \leq \dots \leq x_n^0$ and observe $x_n^0 - x_1^0 \leq 1 - 2\epsilon$. If $x_n^t - x_1^t > \epsilon$, we place a control at $x_1^t + \epsilon < x_n^t$. As an abbreviation we set $R = I_1(x_1^t, \dots, x_n^t)$. At most $n - 1$ of the n voters can be inside the confidence set of voter 1, i.e., $|R| \leq n - 1$, and we have $x_i^t \geq x_1^t$ for all $i \in R$, see Lemma 4.1. With this we conclude

$$x_1^{t+1} = \frac{1}{|R|+1} \cdot \left(x_0^t + \sum_{i \in R} x_i^t \right) \geq \frac{1}{|R|+1} \cdot \left(x_1^t + \epsilon + \sum_{i \in R} x_1^t \right) = x_1^t + \frac{\epsilon}{|R|+1} \geq x_1^t + \frac{\epsilon}{n}.$$

After at most $n \cdot (\lceil \frac{1}{\epsilon} \rceil - 3) \leq \frac{n}{\epsilon}$ stages we can achieve $x_n^t - x_1^t \leq \epsilon$. Then placing the control at $\frac{x_1^t + x_n^t}{2}$ yields the same opinion, which is also inside $[\epsilon, 1 - \epsilon]$, for all voters after at most $\frac{n}{\epsilon} + 1$ stages, i.e., the first step is completed.

In step 2 we proceed as follows. If $x_1^t \in [l, r]$ nothing needs to be done. Due to symmetry we assume $x_1^t < l$ in the following. If $l - x_1^t \geq \frac{\epsilon}{n+1}$ we place a control at $x_1^t + \epsilon$ so that $x_1^{t+1} = x_1^t + \frac{\epsilon}{n+1}$, since all voters influence voter 1. After at most $(n+1)\epsilon$ stages we have $l - x_1^t < \frac{\epsilon}{n+1}$. In such a situation we set the control to $x_1^t + (n+1)(l - x_1^t)$ such that $x_1^{t+1} = l$.

Applying Lemma 4.1 again, we conclude that we can achieve $x_i^t = l$ after at most $\frac{2n+1}{\epsilon} + 2$ for all $i \in I$. Taking the control as x_i^t we can clearly preserve the configuration to later stages. \square

Thus, given enough time (number of stages) we could always achieve the upper bound of $|I|$ convinced voters. By setting $[l, r] = [1 - \epsilon, 1 - \epsilon]$ and $x_i^0 = \epsilon$, we see that the stated estimation gives the right order of magnitude in the worst case.

For the DG-model the upper bound on the time needed to convince all voters depends on the influence w_{i0} of the control for each voter. To this end we define $\omega = \min_{i \in I} w_{i0}$, i.e., the tightest possible lower bound on the influences. Since it may happen that no stable state is reached after a finite number of stages, we can only navigate the voters into an interval of length greater than zero.

Lemma 4.3. *Consider an instance of the DG-model with $0 < \omega = \min_{i \in I} w_{i0} < 1$, $\delta \in (0, 1)$, and a position $p \in [0, 1]$. It is possible to select suitable controls at each stage such that after at most $\frac{\log(\delta)}{\log(1-\omega)}$ stages all $|I|$ voters have an opinion within the interval $[p - \delta, p + \delta]$.*

Proof. By induction over t we prove that we have $|x_i^t - p| \leq (1 - \omega)^t$ for all $i \in I$ and $t \in \mathbb{N}$, if we place the control at position p at all stages. Since $x_i^0, p \in [0, 1]$ we have $|x_i^0 - p| \leq (1 - \omega)^0 = 1$ for all $i \in I$. For $t \geq 1$ we have

$$\begin{aligned} x_i^t &= \sum_{j \in I} w_{ij} \cdot x_j^{t-1} + w_{i0} \cdot p \geq \sum_{j \in I} w_{ij} \cdot (p - (1 - \omega)^{t-1}) + w_{i0} \cdot p \\ &\geq (1 - \omega) \cdot (p - (1 - \omega)^{t-1}) + \omega \cdot p = p - (1 - \omega)^t \end{aligned}$$

for all $i \in I$. Similarly we conclude $x_i^t \leq p + (1 - \omega)^t$. \square

Thus, given enough time (number of stages) we could always achieve the upper bound of $|I|$ convinced voters if the conviction interval has a length greater than zero. By setting $p = 1$ and $x_i^0 = 0$, we see that the stated estimation gives the right order of magnitude in the worst case. Using the Taylor expansion of $\log(1 - \omega)$ and having an influence that decreases with the number of voters in mind, we remark that

$$\frac{\log(\delta)}{\log\left(1 - \frac{1}{n+1}\right)} \leq -(n+1) \cdot \log(\delta).$$

5. COMPUTATIONAL INFORMATION ON OPTIMAL CONTROLS

How can one find reliable information on optimal controls and their resulting optimal achievements? That is, for a special instance like our benchmark instance, how can we find out, how many convinced voters are achievable for a given horizon N ? It is certainly not possible to try all possible controls and pick the best – there are uncountably infinitely many feasible controls, because all elements in $[0, 1]^N$ constitute feasible controls. On the other hand, some logical constraints are immediate without enumeration: it is impossible to achieve more convinced voters than there are voters.

A common technique to supersede such a trivial statement without complete enumeration is to devise a *mathematical model* and find solutions to it by *exact methods*. Exact methods depend on generic logical reasoning or provably correct computational information about the solution of a mathematical model.¹³ In this section, we use *mixed integer linear programming (MILP)* for modeling the DG and the BC optimal control problems and `cplex` for solving them, see Table 11 for the precise parameter settings.

While the models are generically correct, concrete computational results will only be given for our benchmark problem and related data. One big advantage of MILP models is that there is commercial software of-the-shelf that can provide solutions to such models regardless of what they represent. There is even academic free software that is able to provide solutions stunningly fast.

In this section, we will not spell out the formulae of our models explicitly.¹⁴ Instead, we try to emphasize the features of our approach.

First, an optimal solution to an MILP model is *globally* optimal. That is, no better solution exists *anywhere* in the solution space. Second, if an optimal solution to an MILP model was reported by the solver software, we are *sure* (within the bounds of numerical accuracy) that it is an optimal solution, i.e., the method is *exact*. Third, if an optimal solution to an MILP could not be found in reasonable time, then very often we still obtain *bounds* on the value of an (otherwise unknown) optimal solution. And fourth, as usual the process of constructing an MILP model is *non-unique*, i.e., usually there are many, substantially different options to build an MILP model for the same problem, and one may provide solutions faster or for larger instances than another.

We built an MILP model for the DG optimal control problem and two MILP models for the (much more difficult) BC optimal control problem.

5.1. Principle Ideas of the MILP models. The MILP model for the DG optimal control problem can be classified as a *straight-forward* MILP: the system dynamics is linear and fits therefore the modeling paradigm of MILP very well. The only little complication is to model the number of convinced voters, which is a non-linear, non-continuous function of the voters’ opinions. Since binary variables are allowed in MILP, we can construct such functions in many cases using the so-called “Big- M method.” Details are described in the Appendix A.1.

An MILP model for the BC optimal control problem is *not straight-forward* at all. Since the system dynamics depends on whether or not some voter is in the confidence interval of another, we have to decide at some point whether or not two voters’ distance is either less than or equal to ϵ or strictly larger than ϵ . It is another general feature that strict inequalities cause trouble for MILP modeling, and the distinction would be numerically unstable anyway (most MILP solvers use floating-point arithmetic, see Section 3). Thus, we refrained from trying to build a correct MILP model for the BC optimal control problem. Instead, we built two complementary MILP models. Without referring to the details, we again explain only the features. In the first MILP

¹³An example of generic logical reasoning can be seen in the previous section.

¹⁴Mathematically explicit descriptions, suitable for replicating our results, can be found in Appendix A in the appendix.

model, the *lower-bound model*, any feasible solution defines a feasible control, which achieves, when checked with exact arithmetic, at least as many convinced voters as predicted in the MILP model. The second MILP model, the *upper-bound model*, is some kind of *relaxation*: No control can convince more voters than the number of convinced voters predicted for any of its optimal solutions.

This is implemented by using a *safety margin* $\hat{\epsilon}$ for the confidence interval. In the models, it is now required for any feasible control that it leads, at all times, to differences between any pair of voters that are either less than or equal to ϵ or greater than or equal to $\epsilon + \hat{\epsilon}$. If $\hat{\epsilon} > 0$, we obtain a lower-bound model, since some originally feasible controls are excluded because they lead to differences between voters that are too close to the confidence radius ϵ . If $\hat{\epsilon} \leq 0$, we obtain an *upper-bound model* where the requirements for “are within distance ϵ ” and “are at distance at least $\epsilon + \hat{\epsilon}$ ” overlap so that the solution with better objective function value can be chosen by the solver software.

Now, if we put together the information from *both* models then we can achieve more: If the optimal numbers of convinced voters coincide in both models, then we have found the provably optimal achievable number of convinced voters although we had no single model for it. Otherwise, we obtain at least upper and lower bounds. Moreover, any lower bound for the number of convinced voters predicted by the lower-bound model is a lower bound on the number of achievable convinced voters, and any upper bound on the number of convinced voters predicted by the upper-bound model is an upper bound on the number of achievable convinced voters.

The first MILP model for the BC-model is a *basic* model with mostly straight-forward variables along the lines of the DG MILP. However, the system dynamics is discontinuous this time, which requires a much heavier use of the Big- M method. MILP-experience tells us that too much use of the Big- M method leads to difficulties in the solution process. Since the basic model did indeed not scale well to a larger number of rounds, we engineered alternative models.

The resulting *advanced* MILP model for the BC optimal control problem has more complicated variables but not so many Big- M constructions. Moreover, the advanced model uses a perturbed objective function: Its integral part is the predicted number of convinced voters, and one minus its fractional part represents the average distance of the unconvinced voters to the conviction interval. This perturbation was introduced in order to better guide the solution process. The problem with the unperturbed objective function is that many feasible controls unavoidably achieve identical numbers of convinced voters because there are simply much fewer distinct objective function values than controls; this is a *degenerate situation*, which is difficult to handle for the MILP solver. The perturbation guarantees that two distinct controls are very likely to have distinct objective function values.

Our hypothesis was that the advanced model would be easier to solve, which, to a certain extent, turned out to be true. We know of no other method to date that yields more *provable and global* information about optimal controls for the BC optimal control problem.

5.2. Computational Results on the Benchmark Campaign Problem. In the following we report on our computational results. We first compare the effectiveness of the two modeling approaches before we restrict ourselves to the more successful model.¹⁵

The MILP for the DG-model was very effective. It could be solved in seconds for the benchmark problem with homogeneous weights. Eleven convinced voters are possible for any horizon, and, of course, no more. The control is non-crucial here, because homogeneous weights lead to an immediate consensus in the conviction interval. But also for other weights, optimal solutions

¹⁵Because of its special characteristics the advanced model was accepted for the benchmark suite MILIB 2010 (Koch *et al.* 2011). Thus, the new model will automatically receive some attention by the developers of MILP solver software and MILP researchers, which may help to clarify the yet open cases.

can be found for all horizons very fast. The real conclusion is that solving the DG optimal control problem on the scale of our benchmark problem is easy, but there are no mind-blowing observations about optimal controls.

Using the basic MILP revealed that the BC-model is in an all different ball-park. Table 2 shows

# stages	optimal value	CPU time [s]	# variables/binary/integer	# constraints	# non-zeroes
1	3	0.01	2454/ 1738/ 572	3487	11 330
2	4	1.42	4776/ 3377/1122	6809	22 319
3	5	355.33	7098/ 5016/1672	10 131	33 308
4	5–11	3600.00	9420/ 6655/2222	13 453	44 297
5	4–11	3600.00	11 742/ 8294/2772	16 775	55 286
6	5–11	3600.00	14 064/ 9933/3322	20 097	66 275
7	3–11	3600.00	16 386/11 572/3872	23 419	77 264
8	4–11	3600.00	18 708/13 211/4422	26 741	88 253
9	0–11	3600.00	21 030/14 850/4972	30 063	99 242
10	0–11	3600.00	23 352/16 489/5522	33 385	110 231

TABLE 2. Results of the basic MILP model for the benchmark problem with a positive $\hat{\epsilon}$ (provably feasible configurations); number of variables/constraints/non-zeroes for original problem before preprocessing; the time limit was 1h = 3600s; MacBook Pro 2013, 2.6 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM, OS X 10.9.2, `zimpl` 3.3.1, `cplex` 12.5 (Academic Initiative License), branching priorities given according to stage structure.

the computational results for our basic lower-bound model, in particular, with $\hat{\epsilon} = 10^{-5} > 0$ yielding provably feasible controls.¹⁶ In order to really prove that no solutions with a better objective (above the upper bound) exist, we would have to rerun the computations with an $\hat{\epsilon} \leq 0$. We skipped this for the basic model and did this only for the advanced model below, since the information obtained by the advanced model is superior anyway.

The solutions of the basic lower-bound model determine for all stages, via the setting of the control and opinion variables, trajectories of opinions. See Appendix A.1 for a detailed description of all variables. All feasible solutions of our lower-bound model correspond to the data of a BC-dynamics induced by the controls. Thus, these trajectories are exactly the trajectories that would appear in a BC-simulation with the same controls. The numbers of convinced voters induced by the trajectories always correspond to the lower values in the optimal value column of Table 2. The corresponding trajectories of the solutions of our basic lower-bound model are depicted in Figure 2. Here an empty red square represents the opinion of the control at that stage. The circles represent the opinions of the non-strategic agents. The range of influence for each control and the conviction interval are hinted by shaded regions. One can clearly see that the structure of the optimal strategy heavily varies with N . No conspicuous pattern is identifiable – at least for us.

Table 3 shows the computational results for our advanced lower-bound model, in particular, with an $\hat{\epsilon} = 10^{-5} > 0$, i.e., all obtained configurations are feasible, i.e., the given controls provably

¹⁶We have added for each instance the typical information about the MILP problem scale, which is characterized by the number of variables, integrality requirements, and non-zero constraint coefficients in the model. The higher these values the larger and harder the problem is considered in the MILP world. Note that in addition to the obvious decision variables modeling the actual control opinion there is a large number of auxiliary variables necessary to achieve the correct logic. Conventional MILP solvers fall back into a guided enumeration (Branch-and-Bound) of discrete variables whenever they assess that no other information can be computed to their advantage. Branching priorities tell the MILP solver which integral variables should be enumerated first, second, etc.




FIGURE 2. Trajectories produced by solutions of the basic MILP (for $N \geq 4$ we show the incumbent solution when the time limit of 1h was exceeded, i.e., they only show the lower bounds on the optimal values reported in Table 2; for $N \geq 9$ no feasible solution was found in 1h, thus, in those cases we learned no trajectories from the basic model).

# stages	optimal value	CPU time [s]	# variables/binary/integer	# constraints	# non-zeroes
1	3.600	0.02	2367/ 2321/0	6261	62 551
2	4.615	0.18	4656/ 4576/0	11 420	122 736
3	5.640	2.40	6945/ 6831/0	16 579	182 921
4	6.653	113.91	9234/ 9086/0	21 738	243 106
5	6.676	1420.28	11 523/11 341/0	26 897	303 291
6	6.691–11.785	3600.00	13 812/13 596/0	32 056	363 476
7	5.703–11.826	3600.00	16 101/15 851/0	37 215	423 661
8	8.725–11.853	3600.00	18 390/18 106/0	42 374	483 846
9	7.746–11.872	3600.00	20 679/20 361/0	47 533	544 031
10	8.761–11.893	3600.00	22 968/22 616/0	52 692	604 216

TABLE 3. Results of the advanced MILP model for the benchmark problem with an $\hat{\epsilon} = 10^{-5} > 0$ (provably feasible configurations); number of variables/constraints/non-zeroes for original problem before preprocessing; the time limit was 1h = 3600s; MacBook Pro 2013, 2.6 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM, OS X 10.9.2, zimpl 3.3.1, cplex 12.5 (Academic Initiative License), branching priorities given according to stage structure. *Remark:* For a time limit of 24h, 6 stages can be solved to optimality (optimal value 6.697) and for 10 stages we obtain a solution with 11 convinced voters (optimal value 11.800–11.840).

produce this objective function (within the numerical accuracy). For $N = 6$ and larger, `cplex` (ILOG 2014) could not find an upper bound with fewer than 11 convinced voters in one hour. For a time limit of 24h, however, the optimum was determined for $N = 6$, and the optimal number 11 of convinced voters for $N = 10$ could be found. Moreover, a better feasible control

with objective 6.707 for $N = 7$ and of 8.767 for $N = 9$, respectively, could be computed in less than 1h = 3600s on a faster computer¹⁷

The instance for $N = 10$ is special in the sense that the trivial bound of 11 convinced voters is sufficient to prove the “voter”-optimality of a solution with 11 convinced voters. It is yet an open problem to find a configuration that provably maximizes the perturbed objective function of the advanced model. Table 4 shows the computational results for our advanced upper-bound model,

# stages	optimal value	CPU time [s]	# variables/binary/integer	# constraints	# non-zeroes
1	3.600	0.06	2367/ 2321/0	6261	62 551
2	4.615	0.44	4656/ 4576/0	11 420	122 736
3	5.640	7.84	6945/ 6831/0	16 579	182 921
4	6.657	303.81	9234/ 9086/0	21 738	243 106
5	6.668–11.743	3600.00	11 523/11 341/0	26 897	303 291
6	5.697–11.810	3600.00	13 812/13 596/0	32 056	363 476
7	5.701–11.864	3600.00	16 101/15 851/0	37 215	423 661
8	7.729–11.882	3600.00	18 390/18 106/0	42 374	483 846
9	8.741–11.896	3600.00	20 679/20 361/0	47 533	544 031
10	8.762–11.908	3600.00	22 968/22 616/0	52 692	604 216

TABLE 4. Results of the advanced MILP model for the benchmark problem with an $\hat{\epsilon} = -10^{-5} < 0$ (capturing all feasible and possibly some infeasible configurations); number of variables/constraints/non-zeroes for original problem before preprocessing; the time limit was 1h = 3600s; MacBook Pro 2013, 2.6 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM, OS X 10.9.2, `zimpl` 3.3.1, `cplex` 12.5 (Academic Initiative License), branching priorities given according to stage structure.

in particular, with an $\hat{\epsilon} = -10^{-5} < 0$, i.e., the obtained configurations may be infeasible, i.e., the listed objective function values may be different than the results of the true BC-dynamics applied to the computed controls. However, the set-up guarantees that no feasible configurations exist that produce better objective function values. It is apparent that this time not even for $N = 5$ the optimal value of the upper-bound model could be found in 1h. This can be explained: Since in the upper-bound model a solution has more freedom to classify “inside confidence interval: yes or no”, it is harder for the solver to prove that high objective function values are impossible. On a faster computer¹⁸ we obtained for the upper-bound model optimal values of 6.676 for $N = 5$ in 2649s and 6.700 for $N = 6$ in 117 597s, respectively. Thus, we know that more than six voters are neither possible in five nor in six stages.

Table 5 summarizes the results that the advanced model can generate in 1h per computation by combining the results of the lower-bound and upper-bound models. For up to 3 stages the optimal value of an optimal BC-control was found. Regarding the number of achievable voters only (the integral part of the objective function), we know that in 4 stages six voters are possible, but no more. There is, however, a small gap between the optimal values of the lower-bound and the upper-bound model. For 5 and more stages, the upper bound could not pushed below the trivial bound 11 in 1h computation time.

We see that MILP modeling requires a lot of effort. We know, however, of no other method to date that can prove global optimality of a control for $N = 6$ or larger.

¹⁷Mac Pro 2008 2×2.8 GHz Quad-Core Intel Xeon, 21 GB 800 MHz DDR2 RAM, OS X 10.9.5, `zimpl` 3.3.1, `cplex` 12.5 (Academic Initiative License)

¹⁸Mac Pro 2008 2×2.8 GHz Quad-Core Intel Xeon, 21 GB 800 MHz DDR2 RAM, OS X 10.9.5, `zimpl` 3.3.1, `cplex` 12.5 (Academic Initiative License)

# stages	objective of an optimal BC-control	
	after 1h CPU time	all we know from MILP
1	3.600	3.600
2	4.615	4.615
3	5.640	5.640
4	6.653– 6.657	6.653– 6.657
5	6.676–11.743	6.676
6	6.691–11.810	6.691– 6.700
7	5.703–11.864	6.707 –11.864
8	8.725–11.882	8.725–11.882
9	7.746–11.896	8.767 –11.896
10	8.761–11.908	11.800 –11.908

TABLE 5. Summary of the knowledge that the advanced MILP model could collect for the benchmark problem: first, what we know after 1h time limit per computation and second, what we know at all (selected extra computations in **blue**; tightening the upper bound for $7 \leq N \leq 9$ would have been of paramount interest but was out of reach for us, thus no new figures there).

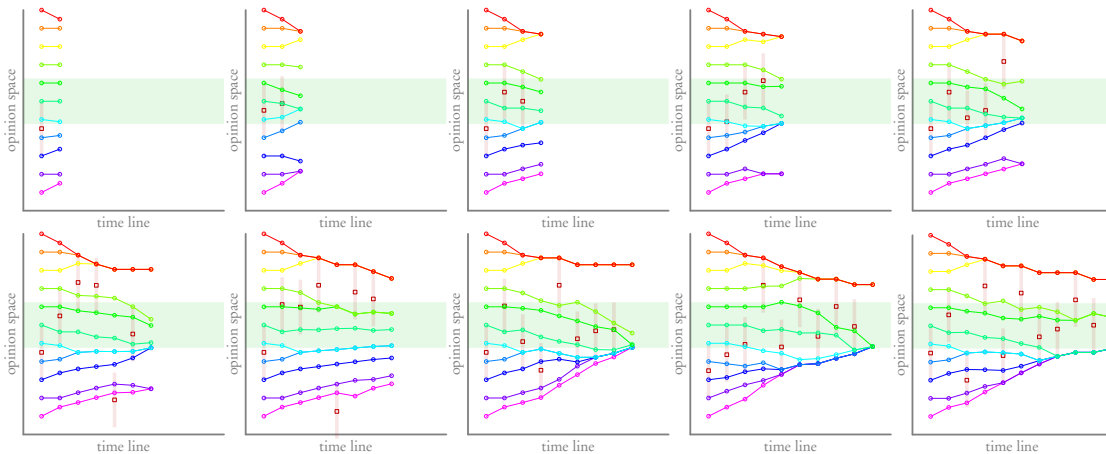


FIGURE 3. Trajectories produced by solutions of the advanced lower-bound MILP for $\hat{\epsilon} = 10^{-5} > 0$ (provably feasible configurations); for $N \geq 6$ we show the incumbent solutions when the time limit of 1h was exceeded

Figures 3 and 4 as well as Figure 5 visualize trajectories. As with the basic model, the optimal solution found by the MILP solver or, in the case of a timeout, the best feasible incumbent solution, determine, via the values of corresponding variables, control values and opinion values for all voters in all stages. See Appendix A.2 for a detailed description of all variables.¹⁹ Again, the numbers of convinced voters induced by the trajectories always correspond to the lower values in the optimal value column of Tab Table 4 for the upper-bound model. The upper values in those tables are produced by involved mathematical computations from MILP solving

¹⁹Note that the voters' opinion values have *not* been taken from the results of a BC-simulation on the basis of the control values; they have rather been taken from the MILP solution vector directly. This way we can better see the differences between solutions to the lower- and the upper-bound model: for the indicated control, the former always produces trajectories that a BC-simulation would have produced, whereas the latter may produce trajectories that are impossible in a BC-simulation.




FIGURE 4. Trajectories produced by solutions of the advanced lower-bound MILP for $\hat{\epsilon} = 10^{-5} > 0$ (provably feasible configurations) for 6 (optimal objective) and 10 (optimal number of convinced voters) stages in a time limit of 24h.




FIGURE 5. Trajectories produced by solutions of the advanced upper-bound MILP for $\hat{\epsilon} = -10^{-5} < 0$ (capturing all feasible and possibly some infeasible configurations); for $N \geq 5$ we show the incumbent solution when the time limit of 1h was exceeded; note the obviously infeasibly splitting trajectories for $N = 10$ stages for voters 1 and 2 in stage 7.

technology (e.g., formal model relaxations). Therefore, they do not correspond to any underlying trajectories.²⁰ Figures 3 and 4 show trajectories induced by the provably feasible solutions of the advanced lower-bound model with $\hat{\epsilon} = 10^{-5} > 0$. Figure 5 shows the (possibly infeasible) solutions of the advanced upper-bound model for $\hat{\epsilon} = -10^{-5} < 0$; infeasibility is best seen for 10 stages where trajectories split, which is impossible for feasible trajectories.

Splitting trajectories in the upper-bound model illustrate the fact that opinion trajectories that are feasible solutions for the upper-bound model²¹ are not necessarily feasible trajectories of opinions under the BC-dynamics. Trajectories coming from the upper-bound model can, e.g., split: For identical opinions i and j whose distance from opinion k is between $\epsilon + \hat{\epsilon} < \epsilon$ and ϵ the optimization algorithm may choose that i is in the confidence interval of k but j is not. This

²⁰Recall that in the summarizing table 5 the *lower* values stem from *lower* bounds (derived from feasible solutions) to the *lower*-bound model and the *upper* values stem from *upper* bounds (derived by MILP machinery) for the *upper*-bound model. Thus, the numbers of convinced voters induced by the trajectories in Figure 5 that stem from *lower* bounds to the *upper*-bound model do not appear at all in Table 5.

²¹Recall, this is characterized by a *negative* safety margin $\hat{\epsilon} < 0$, see Section 5.1 for the idea and A.2 for the full details.

can lead to different opinions of i and j in the upcoming stage. This contrasts the situation in the lower-bound model.²² The lower-bound model is designed in such a way that *all* solutions with opinions at a distance between ϵ and $\epsilon + \hat{\epsilon} > \epsilon$ are infeasible for it.²³

What can we learn from the provably optimal solutions? We see that

- some controls near the end of the horizon just decrease the average distance of non-convinced voters to the conviction interval, i.e., they are induced by the perturbation of the objective function
- the controls of the earlier stages, however, always try to pull a voter as far as possible, i.e., they are exactly at confidence distance to some selected voter

The second observation will be exactly what is used in Section 6 as an idea for a clever heuristics.

5.3. Computational Results on Random Instances. Of course, the choice of the particular benchmark problem is completely arbitrary. We want to provide additional evidence for the observation that DG control is easier than BC control and that the advanced model performs better than the basic one. To this end, we tested all models on random instances. We chose a uniform distribution of start opinions in the unit interval. The conviction interval was constructed by choosing a *party opinion* o uniformly at random in the unit interval and a *conviction distance* Δ uniformly at random between 0.1 and 0.2. The conviction interval then contained all values with a distance at most the conviction distance from the party opinion, cut-off at zero and one, respectively. The confidence distance ϵ was chosen uniformly at random between 0.1 and 0.2 as well. We drew five samples of random data and ran `cplex` (ILOG 2014) with a time limit of one hour on the resulting instances with one up to ten stages. See Table 6 for the data realizations of the five samples used.

Sample	1	2	3	4	5
Start Opinions					
1	0.0001143810805199624	0.02592622792020585	0.0707248803392809	0.1726953250292445	0.05518012076038404
2	0.09233859556083068	0.1850820815155939	0.121328579290148	0.2160895006302953	0.08982103413199563
3	0.1281244478021107	0.3205364363548663	0.2909047436646429	0.5472322533715591	0.2067191540744898
4	0.146755892584742	0.3303348203958791	0.4370619401887669	0.5975562058383497	0.2219931714287943
5	0.2360889763189687	0.4203678016598261	0.5108276010748994	0.609035598255935	0.3637368954633681
6	0.3023325678199372	0.4353223932989226	0.5507979045041831	0.6977288244985344	0.4884111901019726
7	0.417021998534217	0.4359949027271929	0.5693113258037044	0.7148159946582316	0.6117438617189749
8	0.7203244894557456	0.4847490963257731	0.7081478223456413	0.8556209450717133	0.831327840180911
9	0.9325573614175797	0.5496624760678183	0.8399490424990534	0.9006214549067946	0.8707323036786941
10	0.9971848083653452	0.9315408638053435	0.8929469580512835	0.9670298385822749	0.9186109045796587
11	0.9990405156274885	0.9477306110662712	0.8962930913307454	0.9726843540493129	0.9794449978460197
Parameters					
ϵ	0.1387910740307511	0.1698862689477127	0.1040630737561879	0.1224505926767482	0.1354138042860231
o	0.3965807262334462	0.1544266755586552	0.01874801028025057	0.1414641728954073	0.3967366065822394
Δ	0.1186260211324845	0.1204648636096308	0.1040630737561879	0.19762724454960663	0.1765907860306536

TABLE 6. Random data, sampled with 16 significant decimal digits precision.

Table 7 shows the results for an $\hat{\epsilon} = 10^{-5} > 0$; we again skip the upper-bound computation with $\hat{\epsilon} = -10^{-5} < 0$. One very interesting phenomenon can be identified in Sample 5: the number of convincible voters is *not* monotonically increasing with the number of stages available for control. There can be some unavoidable ‘distraction’ caused by other voters. While in two stages we can achieve five convinced voters, in three stages no more than three are possible. This is some more evidence for the assessment that bounded-confidence dynamics can lead to the emergence of various counter-intuitive structures that make the dynamic system appear random and erratic though it actually is deterministic.

²²Recall, this is characterized by a *positive* safety margin $\hat{\epsilon} > 0$.

²³Recall, that this renders all ‘cutting-edge’ solutions infeasible in the lower-bound model although in fact some of them may still be feasible trajectories in the BC-dynamics.

# stages	DG		BC (basic)		BC (advanced)	
	opt. val.	CPU [s]	opt. val.	CPU [s]	opt. val.	CPU [s]
Sample 1						
1	11	<0.01	2	0.01	2.48	0.03
2	11	<0.01	2	1.77	2.49	0.22
3	11	<0.01	2	190.56	2.51	3.62
4	11	0.01	3-11	3600.00	3.49	114.40
5	11	0.01	3-11	3600.00	3.51	1754.67
6	11	0.01	1-11	3600.00	3.52- 8.62	3600.00
7	11	0.01	3-11	3600.00	3.53-11.60	3600.00
8	11	0.01	-∞-11	3600.00	3.54-11.76	3600.00
9	11	0.01	-∞-11	3600.00	3.55-11.78	3600.00
10	11	0.02	-∞-11	3600.00	2.56-11.81	3600.00
Sample 2						
1	0	<0.01	2	0.01	2.72	0.02
2	0	<0.01	2	3.30	2.73	0.18
3	0	<0.01	2	378.88	2.73	2.02
4	0	<0.01	2-11	3600.00	2.73	87.14
5	0	<0.01	1-11	3600.00	2.74	1558.18
6	0	<0.01	0-11	3600.00	2.75- 9.79	3600.00
7	0	<0.01	1-11	3600.00	2.76- 9.83	3600.00
8	0	<0.01	-∞-11	3600.00	2.76-11.85	3600.00
9	0	<0.01	0-11	3600.00	9.80-11.87	3600.00
10	0	<0.01	-∞-11	3600.00	9.81-11.88	3600.00
Sample 3						
1	0	<0.01	2	0.01	2.53	0.01
2	0	<0.01	2	2.04	2.54	0.05
3	0	<0.01	3	233.37	3.54	0.30
4	0	<0.01	2-11	3600.00	3.54	2.30
5	0	<0.01	3-11	3600.00	3.55	53.41
6	0	<0.01	3-11	3600.00	3.55	985.32
7	0	<0.01	3-11	3600.00	3.56- 3.61	3600.00
8	0	<0.01	-∞-11	3600.00	3.56- 8.69	3600.00
9	0	<0.01	3-11	3600.00	3.56- 8.67	3600.00
10	0	<0.01	-∞-11	3600.00	3.57- 8.75	3600.00
Sample 4						
1	0	<0.01	2	0.01	2.49	0.01
2	0	<0.01	2	1.67	2.49	0.05
3	0	<0.01	2	89.49	2.50	0.30
4	0	<0.01	2-11	3600.00	2.51	2.30
5	0	<0.01	2-11	3600.00	2.51	24.45
6	0	<0.01	2-11	3600.00	2.52	278.52
7	0	<0.01	2-11	3600.00	2.53	1281.73
8	0	<0.01	2-11	3600.00	2.54- 7.59	3600.00
9	0	<0.01	-∞-11	3600.00	2.54- 7.67	3600.00
10	0	<0.01	-∞-11	3600.00	2.55-11.71	3600.00
Sample 5						
1	8	<0.01	3	0.01	3.63	0.02
2	11	<0.01	5	1.57	5.64	0.17
3	11	<0.01	3	89.57	3.66	1.51
4	11	0.01	7-11	3600.00	7.67	5.39
5	11	0.01	7-11	3600.00	7.69	36.32
6	11	0.01	3-11	3600.00	7.70	456.57
7	11	0.01	7-11	3600.00	7.72	2590.83
8	11	0.01	7-11	3600.00	7.73-11.79	3600.00
9	11	0.01	7-11	3600.00	7.74-11.83	3600.00
10	11	0.01	-∞-11	3600.00	7.75-11.86	3600.00

TABLE 7. Results of the lower-bound MILP models on random instances; the time limit was 1h = 3600.00s; MacBook Pro 2013, 2.6 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM, OS X 10.9.2, zimpl 3.3.1, cplex 12.5 (Academic Initiative License), branching priorities given according to stage structure.

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 `cplex` (ILOG 2014). In contrast to the DG-model, where we could solve our benchmark problem for any number of stages between 1 and 10 without any difficulty, the instances from the BC-model are harder. Using the MILP-formulation of the previous section we were only able to determine the optimal control up to 6 stages using `cplex`.

The approach using a MILP-formulation has the great advantage that we receive dual, i.e., upper, bounds for the optimal control. For the lower-bound direction (find feasible controls) one can equally well employ heuristics that can determine good controls more efficiently. Note that also the MILP-approach benefits from good feasible solutions, especially if they respect fixed variables in the branch-and-bound search tree.²⁴ So in the next subsections we give three heuristics to find good controls.

6.1. The Strongest-Guy Heuristics. What makes the problem hard, apart from the discontinuous dynamics and numerical instabilities, is the fact, that the control x_0^t is a continuous variable in all stages t . Thus, at first sight the problem is not a finite one. Let us relax our problem a bit by allowing only a finite number of possibilities for x_0^t at any stage and have a closer look at the situation.

By placing a control x_0^t at stage t some voters are influenced by the control while others are not influenced by the control. We notice that the magnitude of influence rises with the distance between the voters opinion x_i^t and the control x_0^t as long as their distance remains below ϵ . So the idea is, even though we do not know what we are doing, we will do it with full strength. Let $c = \frac{l+r}{2}$ be the center of the conviction interval $[l, r]$ and

$$\mu(i, t) = \begin{cases} \min(x_i^t + \epsilon, 1) & \text{if } i \in I \text{ and } x_i^t \leq c, \\ \max(x_i^t - \epsilon, 0) & \text{if } i \in I \text{ and } x_i^t > c, \\ c & \text{if } i = 0 \end{cases}$$

be a mapping from $(I \cup \{0\}) \times \mathbb{N}_{\geq 0}$ to $[0, 1]$. Pulling an agent $i \in I$ at stage t with full strength towards the center of the conviction interval means to place the control at $\mu(i, t)$. The special choice of the function for $i = 0$ attracts the agents towards the center of the conviction interval and aims to avoid oversteering and oscillations. We call this restriction of the problem the strongest-guy heuristics.

Instead of giving the exact values of x_0^t for all stages t we can also give a sequence of indices $[i_0, i_1, \dots, i_{N-1}]$ of those voters that are pulled at maximum strength. The number of such sequences is finite. Therefore, the controls arising this way can be enumerated, in principle, see Algorithm 1.

In Table 8 we give for our benchmark problem the maximum number of convinced voters that can be achieved by using the strongest guy heuristics together with the corresponding index-vector.²⁵ The corresponding trajectories are drawn in Figure 6. We observe that the strongest-guy heuristics improves the best found solution of the ILP approach, given 1h of computation time, for $N = 7$ stages by two additional convinced voters. For $N = 4$ stages the heuristics misses the optimum of 6 convinced voters by one. Given the upper bounds from Table 5 we can conclude that the strongest-guy heuristics found an optimal solution for $N \in \{0, 1, 2, 3, 5, 6, 10\}$.

²⁴A branch-and-bound search tree is built by conventional MILP solvers in order to organize the enumeration of integral variables. In each node of that tree, some of the integral variables have been fixed to a value. A heuristic is only useful in the branch-and-bound solution process if it can find solutions in all nodes of the branch-and-bound tree, i.e., solutions that are consistent with the variable fixings in that node.

²⁵We have used an implementation with exact arithmetic to avoid numerical inaccuracies.

# stages	# convinced voters	control sequence
0	3	[]
1	3	[4]
2	4	[4, 4]
3	5	[3, 8, 0]
4	5	[3, 3, 8, 0]
5	6	[3, 3, 8, 10, 3]
6	6	[3, 8, 8, 2, 7, 2]
7	8	[3, 8, 1, 0, 1, 1, 1]
8	8	[3, 6, 0, 3, 1, 9, 9, 9]
9	8	[3, 8, 6, 0, 9, 9, 9, 9]
10	11	[3, 0, 0, 10, 6, 9, 3, 4, 7, 0]

TABLE 8. Results of the strongest-guy heuristics on the benchmark example.

Algorithm 1: Strongest-Guy Heuristics

Input : ϵ, I, N, x_i^0 for all $i \in I$
Output: best configuration for the strategic agent

```

1 champion  $\leftarrow -1$ 
2 foreach  $[i_0, i_1, \dots, i_{N-1}]$  in  $(I \cup \{0\})^N$  do
3    $\kappa \leftarrow \text{NumberOfConvincedVoters}(\mu(i_0, 0), \mu(i_1, 1), \dots, \mu(i_{N-1}, N-1))$ 
4   if  $\kappa > \textit{champion}$  then                                     /* better solution found? */
5     champion  $\leftarrow \kappa$ 
6     for  $j \leftarrow 0$  to  $N-1$  do                               /* update best solution */
7        $i_j^* \leftarrow i_j$ 
8     end
9   end
10 end
11 return best solution

```

We can improve upon these findings by slightly modifying the strongest-guy heuristics. In order to improve the running time of our implementation, we had replaced the exact arithmetic by floating point numbers. To avoid numerical instabilities we had experimented with the positions $x_i^t + \epsilon - \delta$ and $x_i^t - \epsilon + \delta$ for $\delta = 10^{-6}$, i.e. almost full strength. Curiously enough, we have obtained better solutions in some cases. For $N = 4$ stages the δ -modified control sequence $[3, 3, 8, 6]$, corresponding to the control

$$x_0 = \left(\frac{349999}{1000000}, \frac{309999}{800000}, \frac{550001}{1000000}, \frac{122599789}{200000000} \right) \approx (0.349999, 0.387499, 0.550001, 0.612999),$$

results in six convinced voters. The corresponding trajectory is drawn in Figure 7. We remark that we are not aware of any further improvements.

Having a look at the optimal controls of Table 8 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 $[3, 0, 0, 10, 6, 9, 3, 4, 7, 0]$ or $[3, 8, 8, 2, 7, 2]$?

We remark that the strongest-guy heuristics can be easily adopted to the situation where some of the 0-1 variables from the MILP formulation (see Appendix A) are fixed to either 0 or 1. Thus, it is possible to install a call-back to the strongest-guy heuristics inside an MILP solution





FIGURE 6. Trajectories produced by solutions of the strongest-guy heuristics

FIGURE 7. Trajectories produced by solution of the modified strongest-guy heuristics for $N = 4$

process. However, in our experience the MILP solution needs most of the time to tighten the dual bound.

6.2. A Genetic Algorithm. To get a better idea about the solution space, we implemented a genetic algorithm (GA) to heuristically search for optimal solutions.

GAs work on the analogy of biological evolution and the Darwinian principle of *survival of the fittest*. They test many sets of parameters automatically to find good solutions to the problem. GAs work based on rounds (so called *generations*). Parameter sets resulting in good solutions advance to the next generation, bad ones are discarded. By modifying the good parameter sets before testing them again, new sets are created and tested. Over the course of many rounds the quality of the solutions gradually increases. Eventually one hopes to find a very good solution without having to enumerate by brute-force all possible parameter combinations.

Because a new parameter set is created by modifying existing sets that are already showing useful results, one of the main assumptions of GA is the idea that the solution space is somewhat smooth and that the optimal solution is found close to an already good solution. If the solution

space is heavily fragmented and good solutions are just slightly removed from very bad solutions GAs tend not to work. During the execution of a GA this would manifest in a non-increasing overall solution quality.

6.2.1. *Setup of the GA.* To find a solution for the voting problem, the only parameters that can be influenced are the position of the strategic agent for each of the discrete time events. In GAs each parameter set is called a *chromosome*, a single parameter is a *gene*. In the voter’s game, a chromosome consists of ten genes, which each encode one position of the strategic agent in the opinion space.

To prevent any rounding problems inherent to the `float` and `double` data types in Java (see Section 3), we chose to use fractions instead. The numerator and the denominator are represented as Java’s `BigInteger`. Thus, the calculation is exact and does not suffer from the usual rounding problems that Java encounters with the primitive data types.²⁶

One of the most popular packages to implement GAs in Java is *JGAP – Java Genetic Algorithm Package*²⁷. It provides a framework for the GA that must be extended with the problem specific fitness function and the chromosome implementation. Then it can be used to execute selection, cross over, mutation and evolution on the chromosomes. We used JGAP version 3.3.3.

Each run of the GA start with a *population* of 500 randomly generated chromosomes. Furthermore, we had to determine how many generations the GA should evolve. We set this value to 250 rounds. After 250 rounds, the overall quality of the solutions did not increase, but reached a plateau from which the GA did not find any better solutions.

6.2.2. *Chromosome And Fitness Function.* Each chromosome is a list of the strategic agents position in that particular configuration. To assess the quality of the chromosome, we have to calculate how many of the voting agents are convinced by the configuration after ten rounds of the game. Analogous to the biological *survival of the fittest* GAs evaluate the quality of a solution with a *fitness function*. This function is highly specific to the problem.

To assess the fitness of each chromosome, we started out with the simplest possible fitness function: At the end of the game we counted how many voters were actually convinced. Whether or not a voting agent had been convinced was determined by the party that the agent was closest to. Because we tried to convince voting agents to vote for party 2, any voting agent that was closer to party 2 than to any other party was considered a voter for that party. The parties have fixed positions in the opinion space. They are located at positions 0.25, 0.5, 0.75. Effectively this results in the conviction interval of [0.375, 0.625]. If a voting agent ended up in that interval, it was considered a voter.

This fitness function (called *MaxVoter* (MV)), see Figure 8 for a typical run of the GA with the MV fitness function, suffers several drawbacks. The main one is that it is a discrete function: A voting agent is either convinced or it is not convinced. If there are two chromosomes that both evaluate to *nine* voters one cannot tell whether one of the chromosomes is closer to *ten* voters than the other. For instance, a chromosome *A* could yield the positions

$$A = \{0, 0.38, 0.39, 0.4, 0.45, 0.5, 0.55, 0.6, 0.61, 0.62, 1\},$$

chromosome *B* could yield

$$B = \{0.37, 0.38, 0.39, 0.4, 0.45, 0.5, 0.55, 0.6, 0.61, 0.62, 0.63\}.$$

Both chromosomes result in nine convinced voters. However, chromosome *B* has the two non-convinced voters (at positions 0.37 and 0.63) much closer to the conviction interval than chromosome *A* (non-convinced voters at position 0 and 1). Probably chromosome *B* is closer to an

²⁶See Section 3 for the relevance of this.

²⁷<http://jgap.sf.net>




FIGURE 8. GA with the *MaxVotes* fitness function yielding at most eight voters

optimal solution than chromosome *A*. Therefore it should have a higher fitness value. To account for these shortcomings, we used eight different fitness functions with varying results.

The fitness functions fall into three different categories:

Weighted Sum: This category of fitness functions calculates the weighted sum of all the voters 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 voters on the very edges of the interval to be the most favorable. The idea behind evaluating positions within the interval differently is that voters sitting on the edges of the interval have the greatest effect on voters that are not in the range yet.

Last Remaining: The class of the last remaining fitness functions does not evaluate every voter but restricts itself to convinced voters (counted with weight 1) and the nearest voter that has not reached the conviction interval yet (weighted according to the function). All other voters 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 functions does not evaluate all voters positions but takes into account the distance between the two outmost voters (i.e. the voter with the highest opinion and the voter with the lowest opinion). Because of the order preserving characteristics of the model, these two voters do not change throughout one run, which means, we can just use the distance between the voter 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

original chromosomes										→	new chromosomes									
1	2	3	4	5	6	7	8	9	0	→	1	2	3	4	5	6	G	H	I	J
A	B	C	D	E	F	G	H	I	J	→	A	B	C	D	E	F	7	8	9	0

TABLE 9. Example for a cross over at gene 7. Letters and numbers depict individual genes.

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 voters 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 voters close to 0.5 as possible. *MDBFL2CS* accounts for this and evaluates the positions of the two outmost voters with respect to their distance to the desired midpoint. The two values are added.

6.2.3. *Selection.* At the end of each round, the fitness value for each chromosome is calculated. Based on the fitness value, the core idea of the GA takes place: Those chromosomes with a high fitness value have a larger chance to “survive” and to advance into the next round. Specifically, the GA framework provides two different selection routines. One selection (*weighted roulette selector*, *WRS*) assigns each chromosome a probability of survival based on its fitness value. The population of chromosomes for the next generation is picked according to the probabilities. The other selector (*Best Chromosomes Selector*, *BCS*) discards the chromosomes with the lowest fitness values and lets the other chromosomes survive. This selector converges quicker than the *WRS*. However, its disadvantage is the possibility that it gets stuck in a local optimum, which is not the global optimum.

6.2.4. *Cross Over And Mutation.* Finally, the new population is crossed over and mutated. The *Cross over* operator picks two chromosomes of the new population at random. Then it picks a random number n between 1 and the number of genes in the chromosome (i.e. in the voter’s game a number between 1 and 10). It splits both chromosomes at the n -th gene, resulting in four halves: chromosome 1, part 1 (C1.1), chromosome 1, part 2 (C1.2), chromosome 2, part 1 (C2.1) and chromosome 2, part 2 (C2.2). It “crosses over” the second parts of the chromosomes, resulting in two new chromosomes, made up of C1.1 and C2.2 (chromosome A) and C1.2 and C2.1 (chromosome B). Table 9 shows an example for a cross over between two chromosomes.

The rate how often cross over happens depends on the population size. It is $r = \frac{\text{population size}}{2}$. Every time the operator is invoked cross over happens $\frac{\text{population size}}{r}$. Thus, with a population of 500, cross over happens two times per generation.

The *mutation* operator picks one gene of a random chromosome. It replaces the gene with a random value. The mutation rate was fixed at $\frac{1}{15}$, resulting in 33.33 mutation per generation.

6.2.5. *Evolution.* After each round the GA determines the fitness of each chromosome. Then it selects the fittest chromosomes according to the selector and puts them into the next population. Once the candidate set has been established, the cross over and mutation happen on the candidate set. After both genetic operators have finished, the GA has established the population for the next generation of chromosomes and starts to evaluate them again. This happens until either a satisfactory fitness value has been reached by one chromosome (ie. 11 voters for the voter game) or until the maximum number of generations has been reached (ie. 250). Finally the GA outputs

the best solution. 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 ratio 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.

6.2.6. *Pseudo-Code.* To provide an explanation that is closer to the actual code, the pseudo code for the core GA routines can be found in Algorithm 2.

6.2.7. *Results.* Figure 9 shows 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 voters, 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.




FIGURE 9. Different fitness functions reach different results

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 10 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

Algorithm 2: Genetic Algorithm

Input : populationSize, maxRounds, selector
Output: best configuration for the strategic agent

```

1 population ← InitializeRandomPopulation(populationSize)
2 for  $i \leftarrow 1$  to maxRounds do                                /* fitness evaluation */
3   foreach chromosome in population do
4     EvaluateFitness (chromosome)
5   end
6   if maximum fitness has been reached then                    /* stop evolution */
7     return best solution
8   end
9   newPopulation ← selector(population)                          /* selection */
10  for  $j \leftarrow 1$  to  $\frac{\text{populationSize} \times 2}{1 \times \text{populationSize}}$  do          /* cross over */
11    chromosome 1 ← random chromosome from newPopulation
12    chromosome 2 ← random chromosome from newPopulation
13    remove chromosome 1 and chromosome 2 from newPopulation
14    crossover(chromosome 1, chromosome 2)
15    newPopulation ← add modified chromosome 1 and chromosome 2
16  end
17  for  $j \leftarrow 1$  to  $\frac{1}{15}$  populationSize do                                /* mutation */
18    chromosome ← select random chromosome from newPopulation
19    remove chromosome from newPopulation
20    Mutate(random gene in chromosome)
21    newPopulation ← add mutated chromosome
22  end
23  population ← newPopulation
24 end
25 return best solution

```




FIGURE 10. Performance of three fitness functions with respect to the convinced voters

number of voters. This, of course, poses a problem, since the only way for the GA of evaluating 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 proved 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 11 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 11. 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 voters 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 voters. Figure 12 shows these two runs and the number of voters that each chromosome generated.




FIGURE 12. The two runs for different values for the BCS survival rate

Altogether the GA provides a heuristics 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 voters that the GA could find a strategy for were nine in the example setting. Strategies yielding nine voters 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. The Model Predictive Control Heuristics. The observation that our benchmark problem can be solved to optimality if restricted to just a few stages suggests the following receding-horizon control heuristics (RHC), sometimes also known under the name *model predictive control (MPC)*. See Algorithm 3 for a pseudocode describing MPC.

Algorithm 3: The MPC Heuristics

Input : a confidence radius ϵ ,
a set of agents I ,
a horizon N ,
start opinions \mathbf{x}^0 ,
an MPC horizon $\tilde{N} \in \mathbb{N}$
a solver `OptimalControl`($\bar{\mathbf{x}}$, \tilde{N} , ϵ) for start opinions $\bar{\mathbf{x}}$ on horizon \tilde{N}

Output: a control sequence u^t , $0 \leq t < N$, for the strategic agent

```

1  $\bar{\mathbf{x}} \leftarrow \mathbf{x}^0$                                      /* initialize opinions */
2  $t \leftarrow 0$                                        /* initialize stage counter */
3 while  $t < N$  do                                     /* receding-horizon loop */
4    $\tilde{N} \leftarrow \min(\tilde{N}, N - t)$                  /* cut the MPC horizon to no. of stages left */
5    $(\bar{u}^s)_{0 \leq s < \tilde{N}} \leftarrow \text{OptimalControl}(\bar{\mathbf{x}}, \tilde{N}, \epsilon)$  /* solve MPC-auxiliary problem */
6    $u^t \leftarrow \bar{u}^0$                                /* set next control to first MPC control */
7    $\bar{\mathbf{x}} \leftarrow \text{BC-dynamics}(u^t, \bar{\mathbf{x}}, \epsilon)$  /* update opinions */
8 end
9 return  $u^t$ ,  $0 \leq t < N$ 

```

Although this method uses exact optima of a related optimal control problem, it is in general a suboptimal control. The hope is that its performance is in many cases not too far away from the optimal objective value.

For our computational results we used the advanced MILP to solve the upcoming MPC-auxiliary problems. In order to simplify the setup, we implemented MPC in the following form: Compute the MPC-auxiliary problem, then fix the first control to the solution value, finally increase the number of stages for the MPC-auxiliary problem by one and repeat. This way, the opinions in past stages are computed from scratch in each iteration based on fixed controls. This does not make a difference in the resulting number of convinced voters. Because of the usage of rounded intermediate results (the fixed controls for earlier stages), we observed some infeasible problems. This was cured by changing the control values by $\pm 10^{-6}$, which cured the problem in all cases.

MPC-horizon	# convinced voters	CPU time [s] for last MPC-auxiliary problem
3	6	6.27
4	8	60.41
5	8	163.43

TABLE 10. Results of the MPC heuristics for various MPC-horizons based on the advanced MILP model applied to the benchmark problem with 10 stages; MacBook Pro 2013, 2.6 GHz Intel Core i7, 16 GB 1600 MHz DDR3 RAM, OS X 10.9.2, `zimpl` 3.3.1, `cplex` 12.5 (Academic Initiative License), branching priorities given according to stage structure.

Table 10 shows the computational results for the MPC heuristics. We see that it is unable to find the optimal control with 11 convinced voters in 10 stages. This indicates that the control leading to 11 convinced voters must use control sub-sequences of at least length five that are suboptimal on the shortened horizon. This is another hint that optimal controls exhibit non-trivial structures.

The conclusion of this section on heuristics is that an intelligent tailor-made combinatorial heuristics works best on our very special benchmark problem. The reason for this is most probably that in the fitness landscape of this benchmark problem solutions with more than nine convinced voters are rare and hidden.

For a more general assessment, all methods should compete on a complete benchmark suite with random perturbations of parameters. This, however, goes beyond the purpose of this paper, in which we wanted to advertise the field of optimal opinion control and show some possible directions of future research.

7. INTERPRETATION OF RESULTS

The computational results are twofold: first we have collected results about *optimal controls* including their performance (how many convinced voters are possible?); second we presented results about *models and solution methods* including their effectivity (how tight is the information they compute?) and efficiency (how quickly do they compute that information?).

The results about *models and solution methods* confirm that the BC optimal control problem is much more difficult than the DG optimal control problem. BC optimal control must be modeled with care: the advanced model performs much better in all tests than the basic model. It is the method that, so far, provides the tightest information on the performance of optimal controls in all our benchmark problems. Still, our ability to find provably optimal controls is limited. Using the fast strongest-guy heuristic, we could find lower bounds for the number of achievable convinced voters that are tight for our benchmark problem in all known cases.

The results about *optimal controls* themselves confirm that the BC dynamics causes many interesting effects: Neither is the number of achievable convinced voters monotone in time, nor does every optimal campaign for N stages constitute an optimal campaign for any smaller number of stages. In other words: Enlarging the time horizon in an otherwise unchanged campaign problem requires a completely new solution. Though this sounds plausible in the real-world, we do not claim that this conclusion can safely be transferred to real-world campaign-planning – what we rather state is that it takes no more than the mechanism of BC opinion dynamics to let this effect emerge.

8. CONCLUSION AND OUTLOOK

We have introduced the problem of optimal opinion control by simply allowing one individual to freely choose in each stage its advertised opinion. All efforts to find optimal controls in a small example instance showed that the structure of optimal controls is complicated. Modeling with MILP-techniques is possible, but even sophisticated models are hard to solve. An optimal control for the campaign problem with eleven voters and one through ten stages remains open for seven through nine stages. Ten stages could be solved by the strongest-guy heuristics, which is able to convince all eleven voters in ten stages. Popular meta-heuristics like genetic algorithms and model predictive control could not find this solution.

The fact that the small campaign problem is still largely a mystery, make its investigation interesting for further research. More specifically:

- What happens if the number of voters increases? How do methods to compute optimal controls scale with this parameter? If there are really many voters (like 1000), what is a plausible control concept to gain a significant influence?
- How do optimal controls change with respect to certain properties of the uncontrolled dynamics? Are situations with uncontrolled consensus easier to control optimally? Are situations easier to control optimally where uncontrolled opinions automatically would end up in the conviction interval?

- Are there more counter-intuitive effects? Can it happen that constantly controlling with the central opinion of the conviction interval yields fewer voters than no control at all?
- What happens if we have control restrictions? If each control has to be inside an interval around the true opinion of the controller, how do optimal controls change and are the methods to compute them still applicable? What if the true opinion inducing these restrictions is itself subject to opinion dynamics?

Other directions are the generalization to more than one controller (game theory) and multi-dimensional opinion spaces. That is:

- How should a suitable game be defined in the first place? Should it be one game with one strategy consisting of ten controls? Should it be an iterated game with only one control as a strategy? Are there equilibria in one and/or the other model?
- Do all equilibrium strategies lead to the same ‘election winner’ (= party who ‘convinced’ most voters)?
- Do three or more controllers make a structural difference?
- Is it ‘easier’ (in a sense yet to be defined) to ‘win an election’ with a ‘radical’ (close to zero or one) or a ‘moderate’ (close to 0.5 or in the middle of the competitors) party opinion?
- Does any of these change if the opinion space has a higher dimension?

We think that optimal opinion control opens the door to a wealth of questions posing mathematical problems that are interesting in their own right. Moreover, these questions and also possible answers trigger equally exciting challenges for their interpretations. After all, formally we are investigating only an artificial society, and what we can learn from the results is no output of mathematics – it has to be discussed very carefully.

ACKNOWLEDGEMENTS

We thank the anonymous referees for very valuable suggestions to improve the presentation of the paper.

REFERENCES

- ACEMOGLU, D. & OZDAGLAR, A. (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications* **1**(1), 3–49. URL <http://dx.doi.org/10.1007/s13235-010-0004-1>.
- ALBI, G., HERTY, M. & PARESCHI, L. (2014). Kinetic description of optimal control problems and applications to opinion consensus. *arXiv preprint 1401.7798*, 1–25.
- BHATTACHARYYA, A., BRAVERMAN, M., CHAZELLE, B. & NGUYEN, H. L. (2013). On the convergence of the Hegselmann-Krause system. In: *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM.
- BLONDEL, V. D., HENDRICKX, J. M. & TSITSIKLIS, J. N. (2010). Continuous-time average-preserving opinion dynamics with opinion-dependent communications. *SIAM Journal on Control and Optimization* **48**(8), 5214–5240. URL <http://dx.doi.org/10.1137/090766188>.
- CARLETTI, T., FANELLI, D., GROLLI, S. & GUARINO, A. (2006). How to make an efficient propaganda. *Europhysics Letters* **74**(2), 222–228.
- CASTELLANO, C., FORTUNATO, S. & LORETO, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics* **81**, 591–646. URL <http://link.aps.org/doi/10.1103/RevModPhys.81.591>.
- CHAZELLE, B. (2011). The total s-energy of a multiagent system. *SIAM Journal of Control and Optimization* **49**(4), 1680–1706.
- CHAZELLE, B. (2012). The dynamics of influence systems. In: *FOCS*. IEEE Computer Society.

- DEFFUANT, G., NEAU, D., AMBLARD, F. & WEISBUCH, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems* **3**, 87–98.
- DEGROOT, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association* **69**(345), 118–121.
- DITTMER, J. C. (2001). Consensus formation under bounded confidence. *Nonlinear Analysis* **47**(7), 4615–4621.
- DOUVEN, I. & KELP, C. (2011). Truth approximation, social epistemology, and opinion dynamics. *Erkenntnis* **75**, 271–283.
- DOUVEN, I. & RIEGLER, A. (2010). Extending the Hegselmann-Krause model. *Logic Journal of the IGPL* **18**, 323–335.
- FORTUNATO, S. (2004). The Krause-Hegselmann consensus model with discrete opinions. *International Journal of Modern Physics C* **15**(7), 1021–1029.
- FORTUNATO, S. (2005a). Damage spreading and opinion dynamics on scale-free networks. *Physica A Statistical Mechanics and its Applications* **348**, 683–690.
- FORTUNATO, S. (2005b). On the Consensus Threshold for the Opinion Dynamics of Krause-Hegselmann. *International Journal of Modern Physics C* **16**, 259–270.
- FORTUNATO, S. (2005). The Sznajd consensus model with continuous opinions. *International Journal of Modern Physics C* **16**(1), 17–24.
- FORTUNATO, S., LATORA, V., PLUCHINO, A. & RAPISARDA, A. (2005). Vector opinion dynamics in a bounded confidence consensus model. *International Journal of Modern Physics C* **16**(10), 1535–1551. URL <http://www.worldscientific.com/doi/abs/10.1142/S0129183105008126>.
- FORTUNATO, S. & STAUFFER, D. (2006). Computer simulations of opinions and their reactions to extreme events. In: *Extreme Events in Nature and Society* (ALBEVERIO, S., JENTSCH, V. & KANTZ, H., eds.), The Frontiers Collection. Springer Berlin Heidelberg, pp. 233–257. URL http://dx.doi.org/10.1007/3-540-28611-X_11.
- FRENCH JR, J. R. (1956). A formal theory of social power. *Psychological Review* **63**(3), 181–194.
- GROEBER, P., LORENZ, J. & SCHWEITZER, F. (2014). Dissonance minimization as a micro-foundation of social influence in models of opinion formation. *The Journal of Mathematical Sociology* **38**(3), 147–174.
- HARARY, F. (1959). A criterion for unanimity in french’s theory of social power. In: *Studies in Social Power* (CARTWRIGHT, D., ed.). Ann Arbor: Institute for Social Research: Univer. Michigan, pp. 168–182.
- HEGSELMANN, R. & KRAUSE, U. (2002). Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation* **5**(3). URL <http://jasss.soc.surrey.ac.uk/5/3/2.html>.
- HEGSELMANN, R. & KRAUSE, U. (2006). Truth and cognitive division of labour: First steps towards a computer aided social epistemology. *Journal of Artificial Societies and Social Simulation* **9**(3). URL <http://jasss.soc.surrey.ac.uk/9/3/10.html>.
- ILOG, I. (2014). Cplex 12.5. commercial MILP solver. URL <http://www.ibm.com/software/commerce/optimization/cplex-optimizer/index.html/>.
- KOCH, T. (2004). *Rapid Mathematical Prototyping*. Ph.D. thesis, Technische Universität Berlin.
- KOCH, T., ACHTERBERG, T., ANDERSEN, E., BASTERT, O., BERTHOLD, T., BIXBY, R. E., DANNA, E., GAMRATH, G., GLEIXNER, A. M., HEINZ, S., LODI, A., MITTELMANN, H., RALPHS, T., SALVAGNIN, D., STEFFY, D. E. & WOLTER, K. (2011). MIPLIB 2010. *Mathematical Programming and Computation* **3**(2), 103–163. URL <http://mpc.zib.de/index.php/MPC/article/view/56/28>.
- KRAUSE, U. (1997). Soziale Dynamiken mit vielen Interakteuren. Eine Problemskizze. In: *Modellierung und Simulation von Dynamiken mit vielen interagierenden Akteuren*. (KRAUSE,

- U. & STÖCKLER, M., eds.). Universität Bremen.
- KRAUSE, U. (2000). A discrete nonlinear and non-autonomous model of consensus formation. In: *Communications in Difference Equations* (ELAYDI, S., LADAS, G., POPENDA, J. & RAKOWSKI, J., eds.). Amsterdam: Gordon and Breach, pp. 227–236.
- KRAUSE, U. (2005). 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.
- KRAUSE, U. (2006). Arithmetic-geometric discrete systems. *Journal of Difference Equations and Applications* **12**(2), 229–231.
- KURZ, S. (2012). On the inverse power index problem. *Optimization* **61**(8), 989–1011.
- KURZ, S. (to appear). Optimal control of the freezing time in the Hegselmann–Krause dynamics. *Journal of Difference Equations and Applications*, 1–13.
- KURZ, S. & RAMBAU, J. (2011). On the Hegselmann–Krause conjecture in opinion dynamics. *Journal of Difference Equations and Applications* **17**(6), 859–876. URL <http://www.informaworld.com/10.1080/10236190903443129>.
- LEHRER, K. (1981). Replies. In: *Keith Lehrer* (BOGDAN, R., ed.). Dordrecht: D. Reidel Publ. Co., pp. 223–242.
- LEHRER, K. & WAGNER, C. (1981). *Rational Consensus in Science and Society: A Philosophical and Mathematical Study*, vol. 21. Dordrecht: D. Reidel Publ. Co.
- LIEBRAND, W., NOWAK, A. & HEGSELMANN, R. (eds.) (1998). *Computer Modeling of Social Processes*. Sage Publications.
- LORENZ, J. (2005). A stabilization theorem for dynamics of continuous opinions. *Physica A Statistical Mechanics and its Applications* **355**, 217–223.
- LORENZ, J. (2006). Consensus strikes back in the Hegselmann–Krause model of continuous opinion dynamics under bounded confidence. *Journal of Artificial Societies and Social Simulation* **9**(1). URL <http://jasss.soc.surrey.ac.uk/9/1/8.html>.
- LORENZ, J. (2007). Continuous opinion dynamics under bounded confidence: a survey. *International Journal of Modern Physics C* **18**(12), 1819–1838. URL <http://www.worldscientific.com/doi/abs/10.1142/S0129183107011789>.
- LORENZ, J. (2008). Fostering consensus in multidimensional continuous opinion dynamics under bounded confidence. In: *Managing Complexity: Insights, Concepts, Applications*. Springer, pp. 321–334.
- LORENZ, J. & URBIG, D. (2007). About the power to enforce and prevent consensus by manipulating communication rules. *Advances in Complex Systems* **10**(2), 251–269.
- MALARZ, K. (2006). Truth seekers in opinion dynamics models. *International Journal of Modern Physics C* **17**(10), 1521–1524.
- MIRTABATABAEI, A., JIA, P. & BULLO, F. (2014). Eulerian opinion dynamics with bounded confidence and exogenous inputs. *SIAM Journal on Applied Dynamical Systems* **13**(1), 425–446.
- PINEDA, M., TORAL, R. & HERNÁNDEZ-GARCÍA, E. (2013). The noisy Hegselmann–Krause model for opinion dynamics. *The European Physical Journal B* **86**(12), 1–10.
- POLHILL, G. G., IZQUIERDO, L. R. & GOTTS, N. M. (2005). The ghost in the model (and other effects of floating point arithmetic). *Journal of Artificial Societies and Social Simulation* **8**(1). [Http://jasss.soc.surrey.ac.uk/8/1/5.html](http://jasss.soc.surrey.ac.uk/8/1/5.html).
- RODRIGUES, F. & DA F.COSTA, L. (2005). Surviving opinions in Sznajd models on complex networks. *International Journal of Modern Physics C* **16**(11), 1785–1792.
- STAUFFER, D. (2002). The Sznajd Model of Consensus Building with Limited Persuasion. *International Journal of Modern Physics C* **13**, 315–317.

- STAUFFER, D. (2003). How to Convince Others? Monte Carlo Simulations of the Sznajd Model. In: *AIP Conf. Proc. 690: The Monte Carlo Method in the Physical Sciences* (GUBERNATIS, J. E., ed.).
- STAUFFER, D. (2005). Sociophysics simulations II: opinion dynamics. In: *AIP Conf. Proc. 779: Modeling Cooperative Behavior in the Social Sciences* (GARRIDO, P., MAROO, J. & MUÑOZ, M. A., eds.).
- STAUFFER, D. & SAHIMI, M. (2006). Discrete simulation of the dynamics of spread of extreme opinions in a society. *Physica A Statistical Mechanics and its Applications* **364**, 537–543.
- WEDIN, E. & HEGARTY, P. (2014). The Hegselmann-Krause dynamics for continuous agents and a regular opinion function do not always lead to consensus. *CoRR* **abs/1402.7184**.
- WEISBUCH, G. (2004). Bounded confidence and social networks. *The European Physical Journal B - Condensed Matter and Complex Systems* **38**(2), 339–343.
- WENMACKERS, S., VANPOUCKE, D. E. P. & DOUVEN, I. (2012). Probability of inconsistencies in theory revision. *European Physical Journal B* **85**, 1–15.
- WENMACKERS, S., VANPOUCKE, D. E. P. & DOUVEN, I. (2014). Rationality: A social-epistemology perspective. *Frontiers in psychology* **5**, 1–14. Article 581.
- WONGKAEW, S., CAPONIGRO, M. & BORZÍ, A. (2015). On the control through leadership of the Hegselmann-Krause opinion formation model. *Mathematical Models and Methods in Applied Sciences* **25**(03), 565–585. URL <http://www.worldscientific.com/doi/abs/10.1142/S0218202515400060>.
- XIA, H., WANG, H. & XUAN, Z. (2011). Opinion dynamics: A multidisciplinary review and perspective on future research. *International Journal of Knowledge and Systems Science* **2**(4), 72–91. URL <http://dx.doi.org/10.4018/jkss.2011100106>.
- ZUEV, A. & FEDYANIN, D. (2012). Models of opinion control for agents in social networks. *Automation and Remote Control* **73**(10), 1753–1764. URL <http://dx.doi.org/10.1134/S0005117912100141>.

APPENDIX A. EXACT MATHEMATICAL MODELS FOR OPTIMAL OPINION CONTROL

In this appendix, we present and briefly explain the detailed mathematical Mixed Integer Linear Programming (MILP) models used in this paper. This should allow the interested reader to replicate our results.

The chosen modeling technique is – not surprisingly – much more powerful in the DG-model than in the BC-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 voters 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).

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 BC-model requires some experience in modeling with MILPs, we start with a model for the DG optimal control problem. Later on, when the main principles are explained, we will present a model for the BC-model.

A.1. An MILP Model for the DG Optimal Control Problem. In this section we present a mathematical model, a Mixed Integer Linear Programming model (MILP), for the solution of

the DeGroot optimal opinion control problem. We start with the DeGroot dynamics in order to explain some crucial MILP-modeling techniques in this easier dynamics. These techniques will be used extensively for the bounded-confidence dynamics, in which the logic is considerably more complicated.

The following MILP model is based on standard modeling techniques 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 voters in the various stages; these variables measure the system states. The variables in stage 0 are given as input data (start state / start value).
- For each voter, 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 voters, which can be expressed as follows:

$$\max \sum_{i \in I} z_i. \quad (4)$$

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:

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

So far, we did not restrict the binary variables. A solver would simply set them 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 voter as convinced, then $\ell \leq x_i^N \leq r$ must hold. In other words, the inequalities $\ell \leq x_i^N \leq r$ can be violated when $z_i = 0$, but they must be satisfied whenever $z_i = 1$. Thus, whether or not we demand the restriction $\ell \leq x_i^N \leq r$ depends on the value of a variable. We call such a conditional restriction a *variable-conditioned constraint* and write it as $\ell \leq x_i^N \leq r$ vif $z_i = 1$. The MILP modeling trick can transform such a variable-conditioned constraint into a set of unconditioned constraints in all cases where the violation of the variable-conditioned constraint is bounded.

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

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

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

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

The complete MILP reads as follows:

$$\max \sum_{i \in I} z_i \quad (8)$$

subject to

$$x_i^{t+1} = \sum_{j \in I_0} w_{ij} x_j^t \quad \text{for all } i \in I, t = 0, 1, \dots, N-1, \quad (9)$$

$$\ell - x_i^N \leq \ell(1 - z_i) \quad \text{for all } i \in I, \quad (10)$$

$$x_i^N - r \leq (1 - r)(1 - z_i) \quad \text{for all } i \in I, \quad (11)$$

$$z_i \in \{0, 1\} \quad \text{for all } i \in I, \quad (12)$$

$$x_i^t \in [0, 1] \quad \text{for all } t = 0, 1, \dots, N, i \in I \cup \{0\}. \quad (13)$$

In the following we will not spell out anymore the results of such transformations. Instead, we will present the variable-conditioned constraints literally in order to make the logic more decipherable. The above MILP with literally expressed variable-conditioned constraints reads as follows:

$$\max \sum_{i \in I} z_i \quad (14)$$

subject to

$$x_i^{t+1} = \sum_{j \in I_0} w_{ij} x_j^t \quad \text{for all } i \in I, t = 0, 1, \dots, N-1, \quad (15)$$

vif $z_i = 1$

$$\ell \leq x_i^N \leq r \quad \text{for all } i \in I, \quad (16)$$

end

$$z_i \in \{0, 1\} \quad \text{for all } i \in I, \quad (17)$$

$$x_i^t \in [0, 1] \quad \text{for all } t = 0, 1, \dots, N, i \in I \cup \{0\}. \quad (18)$$

The reader should bear in mind that MILP models with additional variable-conditioned constraints with bounded violation can be transformed into true MILP models. Thus, such models are accessible for standard MILP solvers like `cplex` (ILOG 2014). Modeling languages like `zimpl` (Koch 2004) even support variable-conditioned constraints directly.

The MILP for the DG-model can be solved efficiently by of-the-shelf software like `cplex`. In particular, solving our benchmark problem for any number of stages between 1 and 10 is possible. For example, 11 convinced voters are possible with only one round for uniform weights, and this does not even need the help of a control.

A.2. MILP Models for the BC Optimal Control Problem. The optimal control problem in bounded-confidence dynamics is non-continuous, thus non-linear. Nevertheless, one can construct an MILP model for it by using variable-conditioned constraints in a similar way as in the previous section.

We introduce a real parameter $\hat{\epsilon}$ (meant to be of small absolute value; in our computational experiments we chose $\hat{\epsilon} = \pm 10^{-5}$) with the following meaning: whenever j is not in the confidence interval of i , then $|x_i - x_j| \geq \epsilon + \hat{\epsilon}$ must hold. For $\hat{\epsilon} > 0$, 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 $\hat{\epsilon} > 0$) or to include potentially infeasible solutions (this happens for $\hat{\epsilon} \leq 0$). In the latter case, we grant the optimization algorithm to choose freely in particular whether or not j is in the confidence interval of i whenever $|x_i - x_j| = \epsilon$.

Thus, the model needs to be applied twice: for the identification of feasible solutions we need to set $\hat{\epsilon}$ to something strictly positive, and for the determination of upper bounds on the optimal number of convinced voters we need to set $\hat{\epsilon}$ to at most zero. The larger the absolute value of $\hat{\epsilon}$ is, the more robust the conclusions are against rounding errors.

If we run the model only once with, e.g., $\hat{\epsilon} = 10^{-5}$ we allow for a small inaccuracy in the upper bound obtained by the model. Such an inaccuracy can not be avoided when a standard MILP-solver is used – the most powerful solvers like `cplex` (ILOG 2014), `xpress`, or `gurobi` use bounded-precision floating point arithmetics, and an accuracy of 10^{-6} is a common setting. Any solution that we miss this way, however, would be non-robust in the sense that a slight deviation from the system dynamics would lead to a different objective.

There are several modeling options out of which we present two. The first model extends the MILP for the DG-model using similar techniques to a much larger extent. We could solve it by the standard MILP solver `cplex` up to $N = 4$ in less than an hour. For $N = 5$ the solver could not even get close to a proven optimal solution in weeks. The second model is a carefully engineered, more complicated system comprising some experience in MILP techniques. With the second model we were able to solve the benchmark problem up to $N = 6$ with `cplex`.

We suspect that the solution for $N = 7$ and above requires tailor-made MILP models and solution techniques.

For both our models, we assume that all voters are numbered according to their starting opinion, i.e., $i < j$ implies $x_i^0 \leq x_j^0$ for $i, j \in I$. This saves some work since the order of voters in the opinion space does never change, due to Lemma 4.1.

Our first, basic model uses the following variables:

- The control variables x_0^t , $t = 0, 1, \dots, N - 1$ model the opinions published by the controller, as above. These are the only independent decision variables. The remaining variables are dependent measurements to compute the objective function.
- The state variables x_i^t , $i \in I$, $t = 0, 1, \dots, N$ measure the opinions of agent i in stage t , as above.
- The binary indicator variables $v_{i,j}^t$ are one if and only if voters $i < j$ are within distance ϵ , i.e., they influence each other.
- The binary indicator variables l_i^t , r_i^t , and c_i^t are one if and only if the control in stage t is to the left by a margin of at least $\hat{\epsilon}$, strictly to the right by a margin of at least $\hat{\epsilon}$, or inside the confidence interval of voter i .
- The binary indicator variables z_i , $i \in I$, are one if and only if voter i is within the conviction interval $[\ell, r]$ in the final stage N , as above.
- The measurement variables $\bar{x}_{0,i}^t$, $i \in I$, $t = 0, 1, \dots, N$, denote the contribution of the control opinion x_0^t in the system dynamics formula in stage t ; this variable must equal x_0^t if the control is in the confidence interval of voter i ; it must be zero otherwise.
- The measurement variables $\bar{x}_{j,i}^t$, $i, j \in I$, $t = 0, 1, \dots, N$, denote the contribution of the voter's opinion x_j^t in the system dynamics formula of voter i in stage t ; this variable must equal x_j^t if that opinion is in the confidence interval of voter i ; it must be zero otherwise.
- The count variables k_i^t , $i \in I$, $t = 0, 1, \dots, N - 1$, denote the number of voters in the confidence interval of voter i in stage t .

With this set-up and the aforementioned use of linearized variable-conditioned constraints, we can formulate the following basic model. The logical details are explained right after the

presentation of the MILP.

$$\max \sum_{i \in I} z_i \quad (19)$$

subject to

$$x_i^0 = x_i^{\text{start}} \quad \forall i \in I, \quad (20)$$

$$l_i^t + r_i^t + c_i^t = 1 \quad \forall t = 0, 1, \dots, N-1, i \in I, \quad (21)$$

vif $c_i^t = 1$ then

$$x_0^t - x_i^t \leq \epsilon \quad (22)$$

$$x_i^t - x_0^t \leq \epsilon \quad (23)$$

end

$$\forall t = 0, 1, \dots, N-1, i \in I,$$

vif $r_i^t = 1$ then

$$x_0^t - x_i^t \geq \epsilon + \hat{\epsilon} \quad (24)$$

end

$$\forall t = 0, 1, \dots, N-1, i \in I,$$

vif $l_i^t = 1$ then

$$x_i^t - x_0^t \geq \epsilon + \hat{\epsilon} \quad (25)$$

end

$$\forall t = 0, 1, \dots, N-1, i \in I,$$

vif $v_i^t = 1$ then

$$x_j^t - x_i^t \leq \epsilon \quad (26)$$

else

$$x_j^t - x_i^t \geq \epsilon + \hat{\epsilon} \quad (27)$$

end

$$\forall t = 0, 1, \dots, N-1, i, j \in I: i < j,$$

$$k_i^t = \sum_{j \in I \setminus \{i\}} v_{\min(i,j), \max(i,j)}^t + 1 + c_i^t \quad \forall t = 0, 1, \dots, N-1, i \in I, \quad (28)$$

vif $c_i^t = 1$ then

$$\bar{x}_{0,i}^t = x_0^t \quad (29)$$

else

$$\bar{x}_{0,i}^t = 0 \quad (30)$$

end

$$\forall t = 0, 1, \dots, N-1, i \in I,$$

yif $v_{\min(i,j), \max(i,j)}^t = 1$ then

$$\bar{x}_{j,i}^t = x_j^t \quad (31)$$

else

$$\bar{x}_{j,i}^t = 0 \quad (32)$$

end

$$\forall t = 0, 1, \dots, N-1, i, j \in I: i \neq j,$$

vif $k_i^t = k$ then

$$x_i^{t+1} = \frac{1}{k} \left(\sum_{j \in I \setminus \{i\}} \bar{x}_{j,i}^t + x_i^t + \bar{x}_{0,i}^t \right) \quad (33)$$

end

$$\forall k = 1, 2, \dots, |I| + 1, t = 0, 1, \dots, N-1, i \in I,$$

vif $z_i = 1$ then

$$x_i^N \geq \ell \quad (34)$$

$$x_i^N \leq r \quad (35)$$

end

$$\forall i \in I$$

$$v_{i,j}^t, l_i^t, r_i^t, c_i^t, z_i \in \{0, 1\} \quad \forall t = 0, 1, \dots, N-1, i, j \in I: i < j, \quad (36)$$

$$k_i^t \in \mathbb{N} \quad \forall t = 0, 1, \dots, N-1, i \in I, \quad (37)$$

$$x_i^t \in [0, 1] \quad \forall t = 0, 1, \dots, N, i \in I \cup \{0\}. \quad (38)$$

The objective function (19) counts the number of voters in the conviction interval in stage N . Restriction (20) sets the positions of the opinions in stage 0 to the given start values. Constraint (21) demands (together with the fact that all involved variables are binary) that exactly one of the variables l_i^t, r_i^t, c_i^t must be one. The meaning is that the control is either strictly to the left, to the right, or inside the confidence interval of voter i in each stage t . With restrictions (22) and (23) we request that whenever $c_i^t = 1$ the distance between the control and voter i is no more than ϵ so that the control is really inside i 's confidence interval. In contrast to this, inequalities (24) and (25) make sure that whenever $r_i^t = 1$ resp. $l_i^t = 1$ the control must be to the right resp. to the left with a distance of at least $\epsilon + \hat{\epsilon}$ from voter i so that the control is really outside the i 's confidence interval. Restrictions (26) and (27) make sure in a similar way that the value $v_{i,j}^t$ correctly indicates whether or not i and j are in each others' confidence intervals. The case distinction between a large distance to the left or to the right is unnecessary because of the order of all voters' opinions, reflected by the indices, stays fixed throughout the process. Constraints (28) sets k_i^t to the number of opinions in the confidence interval of voter i . Constraints (29) and (30) compute how much the control's opinion contributes to the next opinion of voter i . This is either the control's opinion in case $c_i^t = 1$ or zero in case $c_i^t = 0$. Similarly, constraints (31) and (32) compute the contribution of voter j to the next opinion of voter i depending on the value of $v_{\min(i,j), \max(i,j)}^t$. Depending on how many opinions are in the confidence interval of voter i , we can now compute its next opinion by restriction (33). Constraints (34) and (35) make sure the classification in variable z_i of being convinced is consistent with the distance of i 's opinion to the conviction interval.

Our second, more sophisticated 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 $c_l, c_r \in \{0, 1\}$, we introduce variables $v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t$ where $v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1$ if and only if the following holds: voter j_{\min} is the minimal index of a voter in the confidence interval of i , voter j_{\max} is the maximal index of a voter in the confidence interval of i , Index $c_l = 1$ if and only if $x_0^t \geq x_i^t - \epsilon$ (i.e., the control is not to the left of the confidence interval of voter i), and Index $c_r = 1$ if and only if $x_0^t \leq x_i^t + \epsilon$ (i.e., the control is not to the right of the confidence interval of voter i). In particular, all variables $v_{i, (j_{\min}, j_{\max}; 0, 0)}^t$ must be zero. The motivation for these variables is that they are indicating the unique *combinatorial confidence configuration* $(j_{\min}, j_{\max}; c_l, c_r)$ of a voter: If $v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1$ then we know by Lemma 4.1 that all voters $j \in I$ with $j_{\min} \leq j \leq j_{\max}$ 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 voter 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 a voter in the conviction interval in stage N , and j_{\max} is the maximal index of a voter 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 voters in stage N is simply $j_{\max} - j_{\min} + 1$.

With the variables above, a logically consistent model can be formulated, which can solve the benchmark instance up to $N = 5$. Some additional engineering effort was required in order to help `cplex` (ILOG 2014) to obtain the optimal value for $N = 6$ as well. For this, we need the following auxiliary variables.

- For each voter $i \in I$ and each stage $t = 1, \dots, N$ we introduce measurement variables λ_i^t and ρ_i^t denoting the left and right distances of voter i to the conviction interval $[\ell, r]$.

The motivation for these variables is that they provide a continuous measurement for how close we are to convince more voters in stage $t+1$. Thus, with these variables we can perturb the objective function to reduce the dual degeneracy of the model, i.e., solutions with identical original objective value up to stage t have distinct perturbed objective values, hinting at which solution has better chances to improve in the later stages.

- For $i, j \in I$ with $i < j$ and $t = 0, 1, \dots, N$ we introduce binary indicator variables $u_{i,j}^t$ with the following meaning: $u_{i,j}^t = 1$ if and only if in stage t the confidence interval of voter j contains voter i . This is the case if and only if in stage t the confidence interval of voter i contains voter j . The motivation for these variables is, first, to transfer the above symmetry relation to a relation among combinatorial confidence configurations and, second, that branching on these new additional variables leads to more balanced subproblems than branching on the variables for the combinatorial confidence configurations.
- In the same spirit, we introduce for $i \in I$ and $t = 0, 1, \dots, N$ binary indicator variables s_i^t with the following meaning: $s_i^t = 1$ if and only if in stage t the control is in the confidence interval of voter i . The motivation is again that a more balanced branching is possible.

The resulting model, presenting all variable-conditioned constraints literally as above, reads as follows. Again, detailed explanations follow the presentation of the model.

$$\max \sum_{(j_{\min} \leq j_{\max})} (j_{\max} - j_{\min} + 1) p_{(j_{\min}, j_{\max})} \quad (39)$$

$$+ 1 - \frac{1}{N} \cdot \sum_{t=1}^N \frac{1}{\ell} \cdot \frac{1}{n} \cdot \sum_{i \in I} \lambda_i^t - \frac{1}{N} \cdot \sum_{t=1}^N \frac{1}{1-r} \cdot \frac{1}{n} \cdot \sum_{i \in I} \rho_i^t \quad (40)$$

subject to

$$x_i^0 = x_i^{\text{start}} \quad \forall i \in I, \quad (41)$$

$$\sum_{\substack{j_{\min} \leq i \leq j_{\max} \\ c_l, c_r \in \{0,1\}}} v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \quad \forall t = 0, 1, \dots, N-1, \quad (42)$$

$$\text{vif } \sum_{\substack{j_{\max} \geq i \\ c_l, c_r \in \{0,1\}}} v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \text{ then} \\ x_i^t - x_{j_{\min}}^t \leq \epsilon \quad \forall t = 0, 1, \dots, N-1, \quad (43)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I, \\ j_{\min} \leq i,$$

$$\text{vif } \sum_{\substack{j_{\min} \leq i \\ c_l, c_r \in \{0,1\}}} v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \text{ then} \\ x_{j_{\max}}^t - x_i^t \leq \epsilon \quad \forall t = 0, 1, \dots, N-1, \quad (44)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I, \\ j_{\max} \geq i,$$

$$\text{vif } \sum_{\substack{j_{\max} \geq i \\ c_l, c_r \in \{0,1\}}} v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \text{ then} \\ x_i^t - x_{j_{\min}-1}^t \geq \epsilon + \hat{\epsilon} \quad \forall t = 0, 1, \dots, N-1, \quad (45)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I, \\ 0 < j_{\min} \leq i,$$

$$\begin{aligned} \text{vif } \sum_{\substack{j_{\min} \leq i \\ c_l, c_r \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \text{ then} \\ x_{j_{\max}+1}^t - x_i^t \geq \epsilon + \hat{\epsilon} \end{aligned} \quad (46)$$

$$\begin{aligned} \text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I, \\ i \leq j_{\max} < n, \end{aligned}$$

$$\begin{aligned} \text{vif } \sum_{\substack{j_{\min} \leq i \leq j_{\max} \\ c_r \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; 1, c_r)}^t = 1 \text{ then} \\ x_i^t - x_0^t \leq \epsilon \end{aligned} \quad (47)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I,$$

$$\begin{aligned} \text{vif } \sum_{\substack{j_{\min} \leq i \leq j_{\max} \\ c_l \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; c_l, 1)}^t = 1 \text{ then} \\ x_0^t - x_i^t \leq \epsilon \end{aligned} \quad (48)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I,$$

$$\begin{aligned} \text{vif } \sum_{\substack{j_{\min} \leq i \leq j_{\max} \\ c_r \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; 0, c_r)}^t = 1 \text{ then} \\ x_i^t - x_0^t \geq \epsilon + \hat{\epsilon} \end{aligned} \quad (49)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I,$$

$$\begin{aligned} \text{vif } \sum_{\substack{j_{\min} \leq i \leq j_{\max} \\ c_l \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; c_l, 0)}^t = 1 \text{ then} \\ x_0^t - x_i^t \geq \epsilon + \hat{\epsilon} \end{aligned} \quad (50)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ i \in I,$$

$$\sum_{j_{\min} \leq j_{\max}} p_{(j_{\min}, j_{\max})} \leq 1 \quad \forall t = 1, \dots, N, \quad (51)$$

$$\begin{aligned} \text{vif } \sum_{j_{\max} \geq j_{\min}} p_{(j_{\min}, j_{\max})} = 1 \text{ then} \\ x_{j_{\min}}^t \geq \ell \end{aligned} \quad (52)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ j_{\min} \in I,$$

$$\begin{aligned} \text{vif } \sum_{j_{\min} \leq j_{\max}} p_{(j_{\min}, j_{\max})} = 1 \text{ then} \\ x_{j_{\min}}^t \leq r \end{aligned} \quad (53)$$

$$\text{end} \quad \forall t = 0, 1, \dots, N-1, \\ j_{\max} \in I,$$

$$\begin{aligned} \text{vif } v_{i,(j_{\min}, j_{\max}; c_l, c_r)}^t = 1 \text{ then} \\ x_i^t = \sum_{j \in I: j_{\min} \leq j \leq j_{\max}} x_j^{t-1} + c_l c_r x_0^{t-1} \end{aligned} \quad (54)$$

end

$$\begin{aligned} \lambda_i^t &\geq \ell - x_i^t & \forall t = 1, \dots, N, \\ & & i \in I, \\ & & j_{\min}, j_{\max} \in I: j_{\min} \leq j_{\max}, \\ & & c_l, c_r \in \{0, 1\}, \\ \lambda_i^t &\geq 0 & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{55}$$

$$\begin{aligned} \rho_i^t &\geq x_i^t - r & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{56}$$

$$\begin{aligned} \rho_i^t &\geq 0 & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{57}$$

$$\begin{aligned} \lambda_i^t &\leq \ell & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{58}$$

$$\begin{aligned} \rho_i^t &\leq r & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{59}$$

$$\begin{aligned} \lambda_i^t &\leq 0 & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{60}$$

vif $p_{(j_{\min}, j_{\max})} = 1$ then

$$\lambda_i^N \leq 0 \tag{61}$$

end

vif $p_{(j_{\min}, j_{\max})} = 1$ then

$$\rho_i^N \leq 0 \tag{62}$$

end

$$\begin{aligned} u_{i,j}^t &= \sum_{\substack{j_{\max} \geq j \\ j_{\min} \leq j_{\max} \\ c_l, c_r \in \{0,1\}}} v_{i,(j_{\min}, j_{\max}; c_l, c_r)}^t & \forall t = 1, \dots, N, \\ & & i, j \in I: i < j, \end{aligned} \tag{63}$$

$$\begin{aligned} u_{i,j}^t &= \sum_{\substack{j_{\min} \leq i \\ j_{\max} \geq j_{\min} \\ c_l, c_r \in \{0,1\}}} v_{j,(j_{\min}, j_{\max}; c_l, c_r)}^t & \forall t = 1, \dots, N, \\ & & i, j \in I: i < j, \end{aligned} \tag{64}$$

$$\begin{aligned} s_i^t &= \sum_{\substack{j_{\min}, j_{\max} \in I: \\ j_{\min} \leq j_{\max}}} v_{j,(j_{\min}, j_{\max}; 1, 1)}^t & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{65}$$

$$x_0^0 \leq \frac{1}{2} \tag{66}$$

$$\begin{aligned} x_i^t - x_i^{t-1} &\leq \frac{n-i+1}{n-i+2} \epsilon & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{67}$$

$$\begin{aligned} x_i^t - x_i^{t-1} &\geq -\frac{i}{i+1} \epsilon & \forall t = 1, \dots, N, \\ & & i \in I, \end{aligned} \tag{68}$$

$$\begin{aligned} x_i^t &\leq x_j^t & \forall t = 1, \dots, N, \\ & & i, j \in I: i < j, \end{aligned} \tag{69}$$

$$\begin{aligned}
x_i^t &\in [0, 1] & \forall t = 1, \dots, N, \\
v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t &\in \{0, 1\} & i \in I \\
& & \forall t = 1, \dots, N, \\
& & j_{\min}, i, j_{\max} \in I: j_{\min} \leq i \leq j_{\max}, \\
& & c_l, c_r \in \{0, 1\}, \\
p_{(j_{\min}, j_{\max})} &\in \{0, 1\} & j_{\min}, j_{\max} \in I: j_{\min} \leq i \leq j_{\max}.
\end{aligned} \tag{70}$$

The main term of the objective (39) determines the number of convinced voters by the help of the variable $p_{(j_{\min}, j_{\max})}$, which is one if and only if j_{\min} is the minimal index and j_{\max} is the maximal index of a convinced voter. The perturbation (40) adds one and subtracts a penalty term less than one from this number. The penalty is essentially the normalized average distance of the non-convinced voters to the conviction interval. The motivation of this perturbation is, that the standard solver, when branching on variables with increasing stage index, has a chance to identify those partial solutions up to a stage that have greater chances (heuristically) to increase the number of convinced voters in future stages. This influences which branches are inspected first and can lead to faster identification of good primal solutions. Restriction (41) fixes the start values, as in the basic model. Constraint (42) demands that exactly one confidence configuration is selected for each voter in each stage. Constraints (43) through (50) makes sure that the selection of confidence configurations is consistent with the opinions and their distances (in an analogous way to the basic model). Constraint (51) models the fact that there can be at most one conviction configuration at the end. If none of the possible conviction configurations is selected then no voter is convinced in the end. Restrictions (52) and (53) make sure that the selected conviction configuration is consistent with the distances of voters to the conviction interval. The dynamics is represented by restriction (54). Note how much simpler the computation of the dynamics becomes with the help of the confidence configuration variables compared to the basic model. So far, the logic of bounded confidence control is complete. The remaining restrictions are heuristic add-ons in order to accelerate the solutions process in a standard solver by means of the additional variables. Constraints (55) through (62) impose bounds on the distances of voters to the conviction interval. If we put them all together, the distance variables are urged to exactly those distances. Constraint (63) and (64) make sure that the additional variables $u_{i,j}^t$ receive values that are consistent with the selected confidence configurations: voters i and j influence each other if and only if one of the confidence configuration variables $v_{i, (j_{\min}, j_{\max}; c_l, c_r)}^t$ and $v_{j, (j_{\min}, j_{\max}; c_l, c_r)}^t$, respectively, for configurations in which i and j influence each other is one. The sum is taken over all such configurations, thus it does not matter which confidence configuration variable contributes the one. Totally analogous is the effect of constraint (65) for the additional variable s_i^t : it is set to one whenever one of the confidence configuration variables of the form $v_{j, (j_{\min}, j_{\max}; 1, 1)}^t$ is one. Some additional cutting planes are provided by restriction (66), which chooses the first control value in the left half of the opinion space. This is possible because the benchmark problem is symmetric. Restrictions (67) and (68) pose bounds on how far an opinion can move in just one stage. Finally, constraint (69) explicitly demands that the order of opinions is consistent with the indices. The remaining constraints (70) through (72) specify the types of the variables.

If one spells out all variable-conditioned constraints in linear restrictions, then one obtains the problem class `rocII` contained in the `MIPLIB 2010` (Koch *et al.* 2011) benchmark suite. The instance `rocII-4-11` (11 voters, 4 stages) is classified as “easy” whereas already `rocII-7-11` is classified as “challenge” (open problem). The full benchmark problem `rocII-10-11` (status “challenge”) is also contained in the suite. The `MIPLIB 2010` suite constitutes the probably most important test bed used by virtually all developers of standard solvers for tuning their software

products, and it may very well be that general MILP research that is totally unrelated to opinion dynamics will lead to the solution of some of our benchmark instances.

APPENDIX B. THE PARAMETER SETTINGS FOR THE MILP SOLVER

simplex tolerance feasibility	1e-09
simplex tolerance optimality	1e-3
mip strategy variableselection	3 (=strong branching)
mip tolerance absmipgap	1e-3
emphasis numerical	yes
timelimit	3600 (in the respective cases)

TABLE 11. The cplex parameter settings that were used for all computations.

Table 11 shows the cplex parameter setting that we used for our computations. This is meant for possible replication of our results. There is no reason to believe that these parameter values are the best possible. They have been set based on our general computational experience in MILP.

RAINER HEGSELMANN, FAKULTÄT FÜR KULTURWISSENSCHAFTEN, UNIVERSITÄT BAYREUTH, GERMANY
E-mail address: rainer.hegselmann@uni-bayreuth.de

SASCHA KURZ, FAKULTÄT FÜR MATHEMATIK, PHYSIK UND INFORMATIK, UNIVERSITÄT BAYREUTH, GERMANY
E-mail address: sascha.kurz@uni-bayreuth.de

JÖRG RAMBAU, FAKULTÄT FÜR MATHEMATIK, PHYSIK UND INFORMATIK, UNIVERSITÄT BAYREUTH, GERMANY
E-mail address: joerg.rambau@uni-bayreuth.de

STEFAN KÖNIG, DATEV E.G., NÜRNBERG, GERMANY
E-mail address: stefan.r.koenig@googlegmail.com

CHRISTOPH NIEMANN, MAIBORNWOLFF GMBH, MÜNCHEN, GERMANY
E-mail address: christoph.niemann@maibornwolff.de