Simulating hidden dynamics

Introducing Agent-Based Models as a tool for linkage analysis

Martin Wettstein

Abstract

Linkage analyses use data from panel surveys and content analyses to assess media effects under field conditions and are able to close the gap between experimental and survey-based media effects research. Results from current studies and simulations indicate, however, that these studies systematically under-estimate real media effects as they aggregate measurement errors and reduce the complexity of media content. In response to these issues, we propose a new method for linkage analysis which applies agent-based simulations to directly assess short-term media effects using empirical data as guideposts. Results from an example study modeling opinion dynamics in the run-up of a Swiss referendum show that this method outperforms traditional regression-based linkage analyses in detail and explanatory power. In spite of the time-consuming modeling and computation process, this approach is a promising tool to study individual media effects under field conditions.

Keywords: linkage analysis, simulation, Agent-Based Modeling

Introduction

One goal of media effects research is to explain macro-social phenomena and developments based on individual media use and effects. To pursue this goal, two distinct perspectives have been taken by scholars of media psychology and public opinion research. As media psychology focuses on the mechanisms behind media effects and studies the effects of arguments, frames, and rhetoric in experimental studies, public opinion research uses

large-scale surveys to observe societal opinion dynamics and link them to media use patterns of recipients. While both fields have been contributing to our understanding of media effects, both the empirical and theoretical integration of individual effects and societal dynamics are challenging. Scheufele (2008) coined the term 'dilemma of media effects research' to describe this situation in which we know much about psychological processes and public opinion but fail to bring these findings together to form one integrated theory or research paradigm.

In recent years, linkage analyses linking panel surveys to parallel analyses of media content have been gaining popularity as a method to study individual media effects at scale (Scharkow & Bachl, 2016). By using media content as independent variables in a quasi-experimental design, these studies explain changes in respondent's answers by the media content they were exposed to; for example to determine the effect of arguments in a campaign (Schemer, Wirth & Matthes, 2012) or to measure framing effects under field conditions (Schuck, Boomgaarden & de Vreese, 2013). Because of their high external validity and high numbers of respondents, linkage analyses are promising tools in studying media effects under field conditions. However, the effect sizes that may be found in these studies are generally small (Scharkow & Bachl, 2016) and the aggregation of weeks of media content to single values that may be used in statistical analyses is far from trivial or consensual.

To amend for these issues, we propose a novel approach to linkage analyses, which does not require *a priori* rules for the aggregation of media content. Instead of aggregating media content, we propose to use agent-based models (ABM) to simulate the immediate effect of media coverage on a panel of respondents between the panel waves. The accumulation of media effects is thereby not defined but may be observed by watching the reactions of individual agents. Unlike other applications of ABM that are used as a proof of concept with randomly generated data or predictive models that simulate various what-if scenarios on which policy decisions may be founded, the approach outlined in this paper uses empirical data both as input and output reference and estimates the model parameters that best describe the unobserved generative processes that lead to changes between panel waves. Comparable to linear models used in regression analyses and analyses of variance, ABMs are thereby applied as explanatory models.

In this paper, we first outline the general approach to using ABM in linkage analyses and introduce an example dataset which has already been used in traditional linkage analyses. Using this data, we illustrate the stepwise implementation of the proposed approach, including the modeling,

estimation, and evaluation of results. Based on the experiences and findings in this pioneering study, we then discuss potential improvements, scope, limits, and methodological challenges of this novel approach in computational communication science.

Background

Linkage Analyses in Media Effects Research

First proposed and implemented by Miller, Goldenberg and Erbring (1979), linkage analyses combine data from surveys and content analyses to study media effects under field conditions. In these studies, respondents are asked for their personally used media outlets, which are then assessed in quantitative content analyses. Based on the individual media use, each respondent is assigned a score for the content (e.g.: bias, tone, or arguments) they were most likely exposed to. The media content is then used as an independent variable in a quasi-experimental field study. This setting allows for the externally valid estimation of the association between media content and individual attitudes, opinions, or actions. If panel surveys are used, this paradigm even allows for causal inference as changes in respondents' answers may be explained by individually used media content between panel waves.

Depending on the time period between panel waves and the number of media outlets the respondents indicate to have used, linkage analyses require large-scale media content analyses. This may be the reason why this appealing approach to media effects research has but scarcely been used until the late 1990ies. Their recent application in several large-scale projects of political communication (e.g.: Kepplinger, Brosius, & Stab, 1991; Peter, 2003; Schuck, et al., 2013; Schemer, et al., 2012), however, have made them more prominent in the past two decades and the availability of automated content analyses renders them feasible even for small projects.

In spite of their growing feasibility and popularity, linkage analyses have two inherent problems that may hamper their success. First, the aggregation of media content in a mean or sum score and the procedure of matching content to respondents are sensitive to measurement errors. Even small errors in self-reported media use and media content analysis are aggregated and multiplied in the procedure of the analysis. As a result, the effect sizes are diminished, as was shown in recent simulation studies (Sharkow & Bachl, 2016). Second, the exact mechanisms by which media effects of consensual or contradictory media content accumulate

over time are largely unknown. Thus, the aggregation of media content to one score per person for later statistical analyses is difficult and hinges on a series of arbitrary decisions by the researcher. These decisions include, but are not limited to, the relative weighting of different media outlets, the treatment of conflicting information, the role of prominence and length of texts, and the relative weighting of texts issued on different days. Regardless of the exact decisions on computing and weighting media content for aggregation, the result is a single value or a select few values which represent the media content of days or weeks. Since this value is the same for all respondents with equal media use and dramatically simplifies the news content between panel waves, the explanatory power of analysis models may be expected to be low, even in the absence of measurement errors.

In reality, media effects are not likely to be caused by a general bias over all stories but are rather composed of a series of individual short-term effects following the perception of each story. Findings from experimental studies indicate that these immediate effects depend on the current attitude of readers, their social environment, psychological traits, and texts they have read before. These complex and time-dependent processes are lost in aggregated linkage analyses. They may, however, be minutely modeled in agent-based simulations.

Agent-based models in linkage analyses

Agent-based models (ABM) are computer programs that simulate the behavior of artificial agents in a virtual environment. These agents may be observed individually or on an aggregate level to assess complex dynamics and emergent phenomena (Gilbert & Troitzsch, 2005; Macy & Willer, 2002). The core components of an ABM are the *agents* that may represent individual or collective actors and have a set of distinct attributes, such as opinions, psychological traits, or propensities for specific actions. These agents are set in a virtual *environment* that may have a specific topography. Agents observe their environment, react to its properties, move around, and interact with other agents they encounter. Their behavior is determined by a set of *rules* that govern how agents see their environment and other agents and how their attributes change.

The simulation is usually segmented in a series of points in time at which the rules are executed. The number of points in time a simulation runs may be predetermined or depend on the emergence of specific phenomena (e.g.: equilibrium, homogeneity, or death of all agents). Since agents are part of a computer program, their state at each point of the simulation

may be assessed and recorded without interfering with the simulation. This allows for a continuous observation of individual agents, interactions, and the complete system.

Previous applications of ABM in social sciences mainly focused on the simulation of complex systems of randomly generated agents that represented members of small communities (Bousquet et al., 1995), comments in online comment threads (Chmiel et al., 2011), or journalists and events in news cycles (Waldherr, 2014). In these applications that seek to explore the boundary conditions and soundness of theories, random data is used to set up the simulation and the result is compared to empirical observations of real-world phenomena to assess the validity of the simulation. In other studies, empirical data is used as an input reference to define a realistic scenario from which one may explore what-if scenarios to predict, for example, consequences of policy decisions (Gilbert et al., 2018; Voinov et al., 2016).

In this paper, we propose to use empirical data both as an input- and output reference of an ABM by simulating the generative process that lies between two empirical observations. Simulating generative processes is not new and is, for example, applied in the Siena algorithm (Snijders, 2017b) that uses stochastic simulations to find probability functions underlying network dynamics. The use of ABM for simulations with empirical input- and output reference, however, is less conventional, except for some notable models explaining the success of political parties (Muis, 2010; Muis & Scholte, 2013).

In the approach we propose for linkage analyses, empirical data collected in panel surveys and content analyses is used to define the initial and final states, as well as the environment of the model. The respondents of a panel survey serve as agents of whom we know both their attitudes at the beginning and the end of the simulation. Their environment is composed of all other respondents and of the daily changing media content that was published between the panel waves. The rules that determine how the agents respond to each other and their media environment at each point in time are derived from theories on short-term media effects.

The only unknown and variable parts of such a model are the parameters that quantify the rules. It remains, for example, to be determined whether an emotional tone increases the arousal of an agent by 0.1 or 0.3 scale points. The values of these parameters are estimated using the next panel wave as a benchmark for the final state of the simulation. In that respect, the estimation of parameters is comparable to other explanatory approaches, such as regression analyses, where the parameters are estimated to explain the data at hand.

When this approach is used to model media effects between panel waves, the parameters directly quantify the individual and immediate effects of media content on respondents. The data used for their estimation, however, are long-term observations of media content and two or more surveys of a panel of respondents that may lie weeks or months apart. Consequentially, the approach of using ABM in linkage analysis may be understood as a tool for estimating short-term media effects and their accumulation, based on long-term observations.

Parameter estimation
The major challenge and most crucial step in this approach lies in the correct and efficient estimation of model parameters. Unlike parameters in linear models, their size may not be estimated directly from the final state. Rather, since agent-based simulations include complex interactions, even small changes in parameters may have unpredictable effects on the final state. Therefore, an estimation process is required that estimates all parameters simultaneously.

One possibility to find the optimal values of parameters is a grid search in which the range of parameters is iteratively narrowed down. For this approach, an initial range of values is specified for each parameter. Within this range, a small number of values are selected for each parameter and the simulation is run for each possible combination of parameter values. The range is then narrowed down to the combinations that yielded the highest scores before repeating the process. Depending on the number of parameters and their interactions, this brute force solution may take some time and computational power. However, it allows for the identification and quantification of local optima and helps to better understand the parameter space.

A more elegant but demanding way of finding parameters is using Bayesian models, such as Monte Carlo Markov Chains (MCMC) or Method of Moments (Koskinen & Snijders, 2007). These models require a set of prior assumptions on covariance structures and the distribution of parameters and errors before applying an iterative Bayesian process to find their optimal posterior distributions. These models are very useful in finding parameter vectors for fully defined functions in which it is possible to infer the value of any single parameter from the values of all others and the data. For example, these models are used in the Siena algorithm (Snijders, 2017a) to find probability functions underlying network dynamics.

Alternatively, numerical optimization algorithms may be used that require very little assumptions on parameter distributions, errors, and

parameter interdependency. Here, genetic algorithms (Charbonneau, 2002) are especially useful because of their simplicity and applicability to various numerical problems. In these evolutionary algorithms, parameters are modeled as attributes of individuals - or genes in chromosomes - that compete and interbreed with each other. The algorithm only requires a quantification for the fitness of a specific combination of parameters (e.g.: Variance explained by the model) to determine which individuals or chromosomes are most fit in this context. Using recombination and mutation on the most appropriate sets of parameters and selecting the fittest in each generation, these algorithms ultimately find best possible combination with respect to the chosen operationalization of fitness.

Depending on previous knowledge about parameter distributions and the complexity of the ABM, different approaches may seem appropriate. All approaches, however, result in the estimation of the single optimal solution and are prone to overfitting the ABM to the sample at hand. Therefore, as in other approaches to estimate non-linear and interdependent parameters for which standard errors of parameters may not be estimated directly, bootstrapping is advised to determine the range within which true parameters most probably lie (Hayes & Scharkow, 2013). Additionally, cross-validation should be used to assess the generalizability of results obtained (Bartlett, Boucheron & Lugosi, 2002).

Example study

We employ the proposed method of using ABM in linkage analyses in a secondary data analysis of an extensive study on opinion dynamics in the run-up of a referendum. The data was collected in 2006 in the run-up of a referendum concerning new restrictions in Asylum law in Switzerland (Wirth et al. 2010). The panel survey included three waves and covered a period of 10 weeks with an initial sample of 1721 participants and a total of 1483 participants finishing all three waves. A parallel content analysis comprised 3907 stories from 22 newspapers and TV news shows in Switzerland which the participants indicated to have used as a source of information during the campaign.

To date, this data has been used in four different original papers using linkage analyses. Matthes (2012) used the data to test the influence of counter-attitudinal media content on the timing of vote decisions. Although no direct media effect was found in this study, an interaction of counter-attitudinal media content and attitude certainty implies media effects for undecided voters. Schemer and colleagues (2012) used the same data to analyze the effects specific arguments in the campaign had on the attitudes

of individuals with authoritarian values. Again, the study mainly found conditional effects which explained an additional 3% of variance in a regression model beyond sociodemographic factors and previous attitudes. Investigating the effect of media framing on individual interpretations, Wettstein (2012) found that media content increased the explained variance of a model predicting issue interpretations by 2% to 4%. Finally, Schemer (2014) found small but significant media effects on racial attitudes, again amounting to increases of 1% to 2% explained variance.

In this paper, we use the same data to predict the effect of media bias and local opinion climate on the attitude toward the referendum. For this purpose, we use the self-reported living place (ZIP code), use of newspapers and TV news shows (dichotomous), attitude toward the upcoming referendum (10-point Likert scale; Wave1: M=6.22; SD=2.86; Wave2: M=6.08; SD=2.98; Wave3: M=6.17; SD=3.01), attitude certainty (5-point Likert scale; Wave1: M=4.11; SD=1.05; Wave2: M=4.15; SD=0.99; Wave3: M=4.27; SD=0.93), media reliance (4 questions, 5-point Likert scale; α=.706; M=5.03; SD=0.85), and political orientation (11-point Likert scale; M=6.02; SD:2.05) of respondents. In the parallel content analysis, all arguments in favor and opposing the referendum were counted for each text. From this data, we computed the bias toward the referendum for each text to it as an independent variable.

A preliminary regression-based linkage analysis, assigning each respondent the mean bias for the media outlets they relied on for political information and the local opinion climate in their living area while controlling for previous attitude, indicated weak effects of both media and social factors. Between the first and second panel wave, only the local opinion climate has a significant effect on attitudes (β= .057, p<.01), and explains an additional 0.5% in explained variance when controlling for previous attitudes. This significant effect persists between the second and third wave (β = .061; p<.01) where a marginally significant effect of media bias (β= .038; p=0.07) is found as well. Both effects, however, only explain 0.5% additional variance when controlling for previous attitude. These small effects are in line with previous linkage analyses and the simulation study by Scharkow and Bachl (2016).

Method

In order to refine this linkage analysis and investigate the effect of social environment and media content in greater detail, we used an ABM to link

media content and survey responses. For the estimation of model parameters, we used a genetic algorithm, as multiple parameters need to be estimated simultaneously and no assumptions on their range and the distribution of errors was possible. The genetic algorithm aimed for maximal explained variance of attitude change between panel waves. To estimate standard errors of the parameter values and their generalizability we employed bootstrapping and cross-validation.

Step 1: Model specification

The ABM was implemented in Python[1] with three main objects interacting in the simulation. The first object represents the population of agents, and is derived directly from the three waves of the panel survey. Each agent is represented as an object within this population by a dictionary of attributes using their responses as values (see Box 1). In addition to the survey responses in all three waves, each agent is assigned geographic coordinates in 100 kilometers[2] based on their ZIP code and one additional attribute for the dynamic attitude toward the referendum. Initially, this dynamic attitude is set to the attitude in the first panel wave. During the simulation, it may be changed according to the rules of the model.

Box 1: *Excerpt from the dictionary defining the agents in the model. Two agents (99003 and 99159) are shown with their attitude certainty (Cert) in three waves, their attitudes (Einstellung_Pro), their media diet, media reliance, sociodemographics, and location.*

```
99003: {
Cert1: 1.0
Cert2: 1.0
Cert3: 'NA'
Einst: 1.03007311516345
Einstellung_Pro1: 1.03007311516345
Einstellung_Pro2: 1.1937762613008
Einstellung_Pro3: 'NA'
Media: [
'24',
'26',
'3',
]
Media_Reliance: 0.866666666666667
age: 34.0
dim_bildung: 0.6
dim_lr: -1.0
dim_spr: 0.0
dim_x: 4.99
```

```
 dim_y: 1.17
 intnum: 99003.0
 sex: 2.0
}
99159: {
 Cert1: 0.8
 Cert2: 1.0
 Cert3: 0.8
 Einst: 0.285994853817719
 Einstellung_Pro1: 0.285994853817719
 Einstellung_Pro2: 0.826484269381761
 Einstellung_Pro3: 0.760015973665868
 Media: [
 '1',
 ]
 Media_Reliance: 0.666666666666667
 age: 19.0
 dim_bildung: 0.6
 dim_lr: 0.4
 dim_spr: 1.0
 dim_x: 6.76
 dim_y: 2.58
 intnum: 99159.0
 sex: 2.0
}
```

The second object in the simulation represents the environment of the agents, which is defined by daily media content. This environment object contains values for each media outlet and each point in time of the simulation (see Box 2). The values represent the number of arguments in favor and opposing the referendum and a measure for media bias, which is calculated by dividing the difference of pro- and contra-arguments by all arguments at this point in time. The bias scores may be accessed by the agents in the course of the simulation and serve to compute their individually observed media bias.

Box 2: *Excerpt from the object representing the media environment of agents. Shown here is the media content of medium #13 on four time points. For each point in time, the number of positive (Arg_Dicho_1), negative (Arg_Dicho_2), and neutral (Arg_ Dicho_9) arguments and the bias for this medium and point in time is provided.*

```
13: {
38878: {
Arg_Dicho_1: 0.5
Arg_Dicho_2: 1.5
```

```
Arg_Dicho_9: 0.0
Bias: -0.5
}
38885: {
Arg_Dicho_1: 0.5
Arg_Dicho_2: 6.0
Arg_Dicho_9: 1.0
Bias: -0.7333333333333333
}
38892: {
Arg_Dicho_1: 4.333333333
Arg_Dicho_2: 7.666666667
Arg_Dicho_9: 0.666666667
Bias: -0.26315789478254853
}
38899: {
Arg_Dicho_1: 8.0
Arg_Dicho_2: 1.0
Arg_Dicho_9: 0.0
Bias: 0.7777777777777778
}
```

The third object is the set of parameters that are used in the rules to compute the attitude change for each agent at each point of the simulation. While the first two objects are static for all simulations, the parameter object is variable and was changed for each run of the simulation by the genetic algorithm.

Two related theories were used to compose a simple set of rules for this ABM. The most basic rule was taken from Social Impact Theory (Latané, 1981) which holds that each individual is constantly influenced by biases in their environment to the extent of aligning themselves with these biases. Drawing on physical models that describe the alignment of magnetic particles, this theory has successfully been applied to models of societal change, such as the diffusion of opinions and stereotypes (Nowak & Latané, 1995; Nowak, Szamrej & Latané, 1990). In addition, we drew upon the Spiral of Silence (Noelle-Neumann, 1974) which states that each individual has a quasi-statistical sense which allows them to constantly monitor the opinion bias in the media and their social surroundings to discern whether they hold a minority or majority position.

Following these considerations, there is one basic rule for attitude change at each point in time which states that the attitude of each agent changes in accordance with the social bias and media bias they encounter,

weighted by their susceptibility to these biases (1). The social bias is the weighted mean of attitudes of all other agents with the weight being the inverted distance to the respective other agent (2). The media bias is the mean bias of all media outlets the agent uses, for each given point in time (3). The susceptibility to social cues is a function of baseline social suscepti- bility and the attitude certainty of the agent (4), whereas the media suscep- tibility is a function of baseline media susceptibility, media reliance, and attitude certainty (5). Finally, the distance used for weighting social bias is a function of geographic distance and ideological distance on the left-right political orientation scale (6).

$$\Delta Attitude = Bias_{Media} \cdot Susc_{Media} + Bias_{Social} \cdot Susc_{Social} \qquad (1)$$

$$Bias_{Media} = \sum_{m}^{Used\ outlets} \frac{Bias_m}{Used\ outlets} \qquad (2)$$

$$Bias_{Social} = \sum_{i}^{N_{Agents}} \frac{Attitude - Attitude_i}{distance_i \cdot N_{Agents}} \sum_{i}^{N_{Agents}} distance_i \quad (3)$$

$$Susc_{Media} = \alpha_1 + \beta_1 Certainty + \beta_2 MediaReliance \qquad (4)$$

$$Susc_{Social} = \alpha_2 + \beta_1 Certainty \qquad (5)$$

$$distance_i = \sqrt{(X - X_i)^2 + (Y - Y_i)^2} + \beta_3 \sqrt{(Ori - Ori_i)^2} \quad (6)$$

Although this set of rules is quite simplistic in assuming that all respon- dents just align their own opinions to some degree to their social and media environment, there are five parameters which have to be estimated: The baseline media (α_1) and social (α_2) susceptibility, the influences of attitude certainty (β_1) and media reliance (β_2), as well as the relative importance of geographic and ideological distance (β_3). These parameters are used in a simulation over a fixed amount of points in time that represent the time between the survey waves. As we do not have any reason to believe that either of these parameters remains constant in the run-up of a referendum campaign, we do not impose any equality constraint and estimate the para- meters in each panel interval independently.

To put it in other words: In this simplified simulation of the generative process that may lead to attitude changes between panel waves, each agent accurately observes their environment at each point in time. They compute the bias of pro versus contra messages in the media outlets they use and

they get an accurate impression of the opinion climate in their environment to which people living close to them and people with similar political orientation contribute more strongly than others. The degree to which someone has to live apart from an agent to be considered as distant as a person with a different political orientation is an unknown parameter (β_3). Depending on a general media susceptibility of all respondents (α_1) and a general susceptibility to social pressure (α_2), these observed biases have some small effect on the individual respondent. Since not all respondents are equally susceptible, media susceptibility is increased for respondents with high media reliance to some degree (β_2), just as both media and social influence are reduced by some degree (β_1) for respondents that reported high attitude certainty.

Step 2: Internal validity assessment
Before estimating parameters or applying a model in research, an assessment of its appropriateness and internal validity is advised (Waldherr & Wettstein, 2019). In this first test, the implementation of the rules, the behavior of the simulation and the sensitivity to changes within the model are explored. To achieve this, we run the simulation with manually set parameters to observe the emerging opinion dynamics between panel waves. If the model behaves as would be expected with the given environment, starting point and parameters, we may proceed to parameter estimation.

In order to test the implementation of rules and the sensitivity to parameter adjustments, we use unrealistic parameter values that allow only for social influence, only for media influence, and for both influences. The conditions of the referendum campaign under investigations were special in that most of the respondents were in favor of the referendum while most of the media outlets published arguments against it. We therefore expect the model with strong social influence to result in a final state where all people are uniformly in favor. In the model with strong media influence but no social influence we expect most media users to be against the referendum in the final state, while respondents that indicated not using any media outlet to stick to their opinion. If both media and social influence are admitted, even media abstinent respondents will become opposed to the referendum, as they receive second-hand media content through their peers.

The results of this initial test (see Figure 1) are in line with expectations. Since the figure only depicts the mean opinion climate in each 10km by 10km region, individual agents may not be observed. However, there are some squares that do not change their mean attitude as a result of media
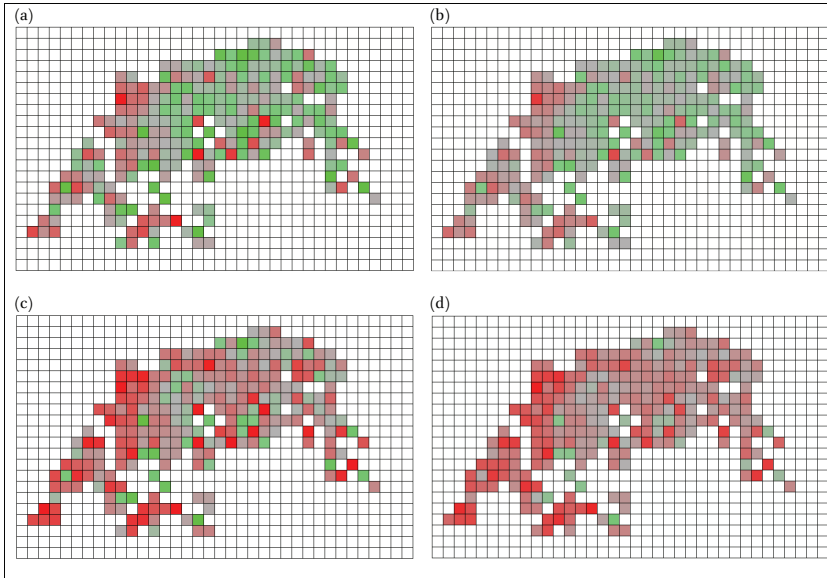
Figure 1. Simulated Opinion climates at panel wave 2 using different parameter sets: a) Initial situation at wave 1; b) Strong social influence; c) Strong media influence in the absence of social influence; d) both influences strong.

Note: The survey was limited to the German and French part of Switzerland. Due to the lack of respondents, the Italian speaking south is not visible on this grid, as are some remote regions in the Alps.

content (Figure 2c) because the respondents living there did not report using any of the analyzed media outlets. Accordingly, some squares retaining a strong pro attitude in the face of contra media bias are located in the south and in central Switzerland, where we do not have data on all regional newspapers.

Furthermore, the test shows that the opinion climate gets blurred once a social influence is enabled. This finding is in agreement with expectations as the mutual influence of agents reduces extreme opinions and leads to a more balanced opinion climate. The initial difference in opinion climate between eastern and western regions of Switzerland remains intact in all models.

As a second test of internal validity, the individual careers of agents may be observed to determine whether any of them behave in a way that would be inconsistent with expectations. As we are investigating the run-up of a referendum, we do, for example, not expect people to fundamentally change their attitudes each day. Likewise, the change of attitude is not likely to

be linear but follows the media coverage and may change according to the environment of each agent. Here, we also use a visual test in which only twelve randomly selected individuals are traced throughout the period under investigation.

The results depicted in figure 2 use the same three different parameter sets as in the previous visualization and the same agents in each example. As expected, there is a convergence to the mean in the condition with only social influence as the extreme opinions are adjusted to their social environment. In the condition with pure media influence, parallel agent careers indicate respondents with the same media diet; the two horizontal lines are respondents that did not indicate to have used any media outlet. In the condition without any influence, the attitudes remain stable, whereas the condition with both media and social influence combines the negative trend caused by the media and the fanning-in caused by social adjustment.

This first validation of the ABM indicates that the rules are specified correctly, the agents have access to their media content, and the model behaves as expected. However, the parameter values used in this test were arbitrary and none of the combinations led to the opinion climate observed
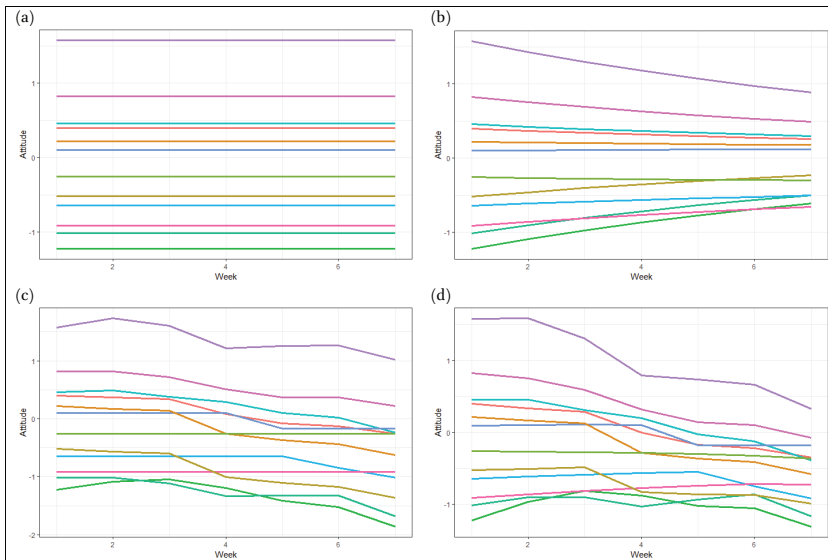

Figure 2. Individual agent careers using different parameter sets: a) All parameters set to zero 1; b) Strong social influence; c) Strong media influence in the absence of social influence; d) both influences strong.

in the second panel wave. The next step, therefore, is to estimate parameter values that lead to a final state of the model close to the observations in the second panel wave.

Step 3: Model parameter estimation
In order to estimate the model parameters, a genetic algorithm was designed and implemented in Python. The algorithm was informed with prior distributions of all five parameters as initial gene pool. The prior distributions were based on the null hypothesis and had a mean value of zero and a small standard deviation. An exception was the relative weight of geographic and ideological distance, which was initially set to 1, indicating no preference. Preliminary tests with sub-samples were used to estimate reasonable values for the prior standard deviation of each parameter.

The genetic algorithm was fashioned to mimic an evolving population of parameter sets that compete and breed to generate new sets following a fixed routine: Initially, the algorithm loads the prior distributions for all parameters and randomly draws parameter values from these distributions to generate 30 complete parameter sets with random values for all five parameters. For each of these parameter sets, the simulation is run once, each time comparing the final attitude of each respondent with their reported attitude in the next panel wave. Agreement with the empirical benchmark is quantified as the inverted mean square deviation of simulated and measured attitudes. The score of agreement is stored for each parameter set before proceeding.

After a complete generation has been tested in this way, a fraction of parameter sets (e.g.: 40%) with the lowest agreement are eliminated. The remaining sets are then used to compute a posterior distribution for each parameter. From these posterior distributions, new parameter sets are randomly generated to replace the eliminated sets. To account for random mutations, there is a small chance (e.g. 10%) that the standard deviation of the posterior distribution is doubled for one random draw of new parameter values. The share of parameter sets that are killed off in each generation and the probability of mutation events may be chosen by the researcher. Higher mutation rates lead to increasingly widespread distributions in parameters with low selective pressure (i.e.: parameters that have little effect on model fitness) whereas lower mutation rates reduce the efficiency of the genetic algorithm as they may lead to quick convergence on suboptimal values or local optima. Likewise, eliminating higher shares of unfit individuals (i.e. parameter sets) in each generation leads to faster convergence at

the risk of finding local optima. Preliminary tests are advised at this stage to determine sensible settings and assess the impact of different decisions.

With only the most optimal parameter sets surviving and producing offspring, the genetic algorithm is bound to converge on optimal combinations of parameter values. When this happens, both the variance of the posterior distribution of parameters and the variance of fitness-scores are reduced. Consequently, the variance of parameter values and fitness-scores may be used as a measure of convergence to decide when to stop the evolutionary process. Theoretically, it would be possible to let the simulation run until the standard deviation of results is exactly zero as the optimal value for each parameter is found up to the maximal precision of the variables and the values are equal for all individuals. Considering the problem of overfitting the parameters to the sample and the number of generations it would take, this does not make sense, however. Alternatively, a standard deviation of results below 10^{-5} or a mean standard deviation of parameter values below 10^{-3} may be set as benchmarks to terminate the algorithm. Again, preliminary tests and visual confirmation of convergence is advised.

For the example study, the algorithm was set up with a selection share of 40% and a mutation rate of 10%. As a measure of convergence, we used the decimal logarithm of the standard deviation of the agreement between the final state and the empirical benchmark. The convergence of parameters in the genetic algorithm over the first 300 generations is shown in figure 3 to illustrate the progression of the genetic algorithm. As may be seen from this illustration, the convergence of parameters started at different times for different parameters. While both the intercept of social impact and the effect of attitude certainty approach their optimal values within 50 generations, the convergence of media impact and the effect of media reliance are much slower. The value of the fifth parameter, the relative weight of geographic distance and ideological mismatch in calculating agent distances, is the last to begin movement but quickly finds its optimal value. The rapid fall of the convergence measure indicates that the model reached a state where the standard deviation of fitness scores differed by less than 0.0001 within 150 generations. At this point, however, the values of four of the parameters were still moving and the variance of three parameters (α_1, β_2, β_3) was higher than the variance of results. From this observation, we may conclude that the exact parameter values of media impact and the relative distance only have a marginal effect on the overall explanatory power of the model.
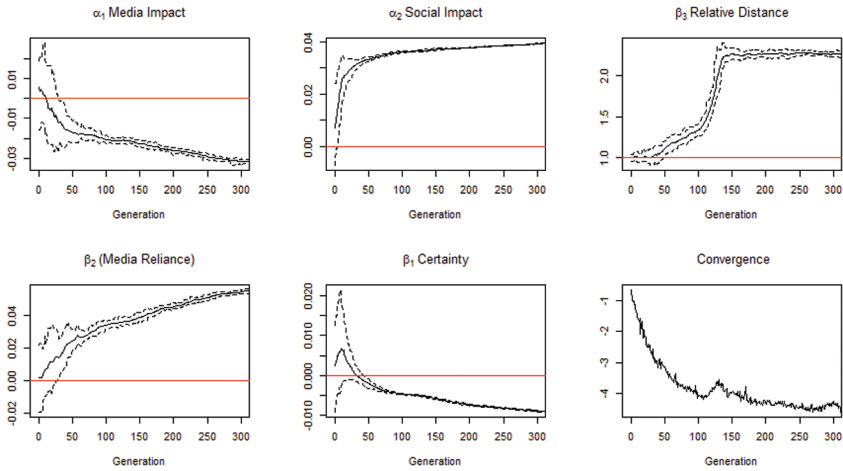
Figure 3. Convergence of parameters and agreement with the benchmark in the genetic algorithm over first 300 generations.

Notes: Dashed lines describe the 95% confidence interval of parameter values in posterior distributions. Horizontal lines represent the null hypothesis and prior means of parameters. For model convergence, the log10 of the standard deviation of agreement with the benchmark is used.

Step 4: Cross-Validation and bootstrapping
Genetic algorithms quickly converge on optimal values and the small standard deviations of parameter values may lead to the illusion that these are the true values with only a small margin of error. This is not correct, however, since the genetic algorithm only finds the optimal solution for the sample that is provided. Whether the parameter values would also be sensible choices for other samples of the same grand total may not be inferred from their posterior distributions.

In order to check for the validity of results, it is therefore advised to cross-validate the results using a different dataset from the same source. Unless the parameters estimated from a sub-sample may be used to model the behavior of other sub-samples, the results may not be considered generalizable. We therefore advise to split the dataset to two different sub-samples, as has become the standard approach in cross-validation in machine learning (Bartlett et al., 2002). One sample serves as a training dataset to find optimal parameters, while the other is used as a test sample to determine how well the parameters are applicable to other samples from the same source.

Furthermore, because the result of the genetic algorithm is one optimal set of parameters specifically fitted to the sample, it is not possible to compute their standard errors. In order to determine the range within which

the true parameter values lie for any sample, we propose to use bootstrapping to estimate probability distributions. That is, the parameters are estimated for different samples, each drawn from the original sample, allowing for repetition (Efron, 1981). While bootstrapping is still dependent on the original sample and the mean optimal parameters are not likely to change, this approach allows for the estimation of probability distributions and reasonable confidence intervals for each parameter.

In essence, there are two possible ways of bootstrapping to obtain probability distributions in genetic algorithms. First, you may use the algorithm to find the single optimal solution, which is then used as prior in *post-hoc bootstrapping*. In this approach, the optimal values are used as new priors for the genetic algorithm which is used to re-estimate the parameters on a small number of bootstrapping samples (N=20-50) of respondents. The variance of optimal solutions for different samples is then used to estimate standard errors and reasonable upper and lower bounds for the true value of each parameter.

Alternatively, *integrated bootstrapping* may be used during the original parameter estimation. In this approach, the genetic algorithm uses a new bootstrapping sample for each generation. The parameters are thereby not fitted to a single complete sample of respondents but on a multitude of constantly changing bootstrapping samples. Consequentially, the parameters are not narrowed down to single optimal values but eventually converge to stable posterior distributions which lead to valid results for all bootstrapping samples.

Both approaches may be used to find posterior parameter distributions that are not overfitted to the sample at hand but allow for the estimation of confidence intervals for true parameter values. Since both approaches use bootstrapping samples, their results are bound to reach similar values.

For the example study, a sub-sample of 1000 cases was used to estimate parameter values before cross-validating them on the remaining 721 cases. After cross-validation, all cases were used to complete the parameter estimation before applying post-hoc bootstrapping to determine their standard errors.

Results

For the estimation of the first interval between panel waves 1 and 2, a sample of 1000 respondents was drawn randomly to estimate parameter values. The genetic algorithm was run for 2000 generations until the standard

deviations of all parameters were close to zero. They ranged from 5.48×10^{-8} for parameter β_2 to 2.80×10^{-4} for parameter β_3. The parameter values were then cross-validated on the 721 remaining respondents. In the training sample, the ABM using the optimal parameter set for this sample was able to explain 8.66% of attitude variance between the panel waves. Applying the model to the test cases that were not included in the training, 10.95% of attitude variance could be explained. While this result seems puzzling at first, it may be explained by the fact that the upper limit of explained variance strongly depends on the composition of the sample, namely the attitude variance in the sample, the number of media users, and the geographical distribution. However, the results indicate that parameter estimations from sub-samples were generalizable for all cases.

Next, the estimation of optimal parameters was extended to the complete sample of 1721 cases to find the initial values of parameters for bootstrapping. The *post-hoc* approach for bootstrapping was used with the optimal parameter values as prior means and the standard deviations large enough to include zero in the 95% confidence interval. Using these priors, 50 samples were drawn and for each the genetic algorithm was run for 70 generations. The number of generations was found to be adequate as the results for all samples converged to standard deviations below 0.0001. Figure 4 depicts the parameter convergence for all 50 samples. The final
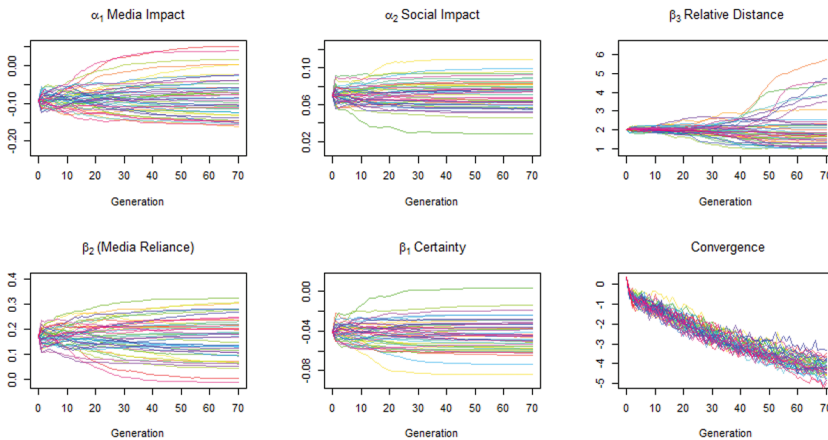

Figure 4.   Post-hoc boostrapping of all parameters.

Notes: The lines indicate the moving average of parameters when adapting to bootstrapping samples for 50 generations. Convergence is measured as the logarithm of the standard deviation of agreement. For all bootstrapping samples, this standard deviation dropped below 10^{-3}.

values are distributed around the initial value, as would be expected in bootstrapping. From the distributions, we computed the standard errors and posterior distributions for all parameters.

The complete procedure was then repeated for the second interval between the second and third panel wave. Again, 50 bootstrapping samples and 70 generations were used to determine the standard errors for all parameters. The ABM using optimal parameters for each interval was able to explain 9.8% (Wave 1-2) and 8.1% (Wave 2-3) of variance of attitude change and the confidence interval of most parameters after bootstrapping did not include zero (Figure 5). Thus, the explanatory power and interpretability substantially exceeded the regression-based linkage analysis reported above. In addition, since moderating effects are directly modeled in the rules and are not expressed by multiplicative terms in the model, the interpretation of conditional effects in ABM is intuitively interpretable.
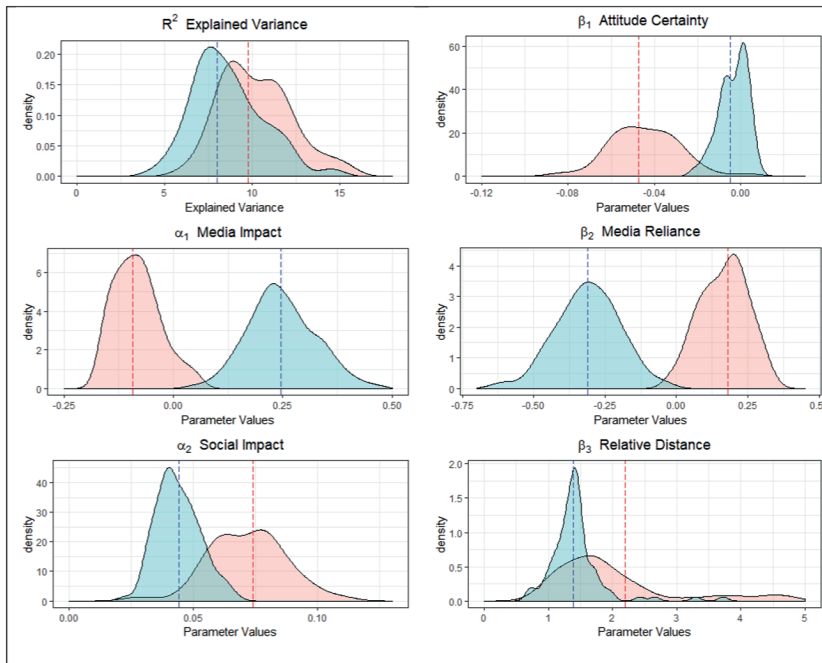

Figure 5. Density plots of explained variance and optimal parameters after post-hoc bootstrapping.

Notes: Red: Estimations for the first panel interval; Blue: Estimations for the second panel interval. Dashed lines indicate the point which was found to be the optimal parameter value for all 1721 cases in either interval.

The results (see Figure 5) suggest that the media and social environment of respondents had a considerable impact on their attitude toward the referendum. Specifically, we found that the social bias had a significantly positive effect in both panel intervals and was moderated by attitude certainty only between waves 1 and 2. In the second interval, attitude certainty had no effect on social or media influence. Media influence was positive only for persons with high media reliance in the first interval, whereas it was generally positive between panel waves 2 and 3, except for respondents with high media reliance. For the relation of geographical and ideological distance, we found similar values for both panel intervals. The coefficients indicate that an ideological distance of one point on the 11-point left-right scale was equal to 50-80 km in the first interval and 70-80 km in the second, indicating a slight ideological polarization in the run-up of this campaign.

With regard to ABM as an explanatory approach in linkage analyses, the results of this example study are encouraging. It was possible to pin-point media effects during a referendum campaign, their dependency from media reliance, and their relation to social influences. Using rules to explicitly model baseline and conditional effects, the moderating effect of attitude certainty and media reliance could be estimated directly. Here, the model outperforms traditional regression models that rely on multiplicative interaction terms to determine moderating effects at the cost of multi-collinearity and distorted intercepts. Also, the direct estimation of moderating effects is intuitively interpretable and does not depend on the size and direction of main effects.

Running the model under various conditions, we found that the explanatory power of the ABM was slightly increased when panel drop-outs were included in the simulations. Even if the career of these agents may not be compared to empirical measurements in later panel waves, they contribute to the social environment of other agents and may relay media effects to their peers. Other than in regression-based models for panel data analysis, these cases are not lost to the analysis but contribute valuable information to the model.

With regard to evolutionary parameter estimation, we found that this approach was easy to implement and yields satisfactory results but is extremely time-consuming. While each single simulation of the whole panel just takes a few seconds, one generation with 30 parameter sets took about 50-70 seconds to compute and assess[3]. As the genetic algorithm required more than 3000 generations to reach optimal parameters in the final run and 3500 generations to determine the standard errors, the estimation of parameters took several days. Counting preliminary analyses required for

sensible priors and additional analyses using cross-validation and performance tests, more than one week was required to compute the results for this study. Furthermore, as agent-based modeling is an iterative process in which rules are continually revised, included, or removed and each model takes hours to estimate, the computer was calculating simulations for eight weeks straight for this paper. In this respect, the approach of using ABM in linkage analyses is clearly outperformed by aggregate linkage analyses using regression models which take few hours to prepare and seconds to compute.

Methodological challenges and solutions

Taken together, the results presented for this example study are promising. They indicate that using ABM in linkage analyses does not only increase the explanatory power but allows for more complex and detailed modeling of media effects. In spite of its appeal, however, there are some challenges to this approach that need to be addressed. Most importantly, researchers need to keep in mind the challenges in model development, the efficient parameter estimation, and overfitting, which we elaborate in this section.

Model development
One crucial challenge in all applications of ABM lies in the model specification. There is no upper limit to model complexity and it is generally tempting to add rules and conditions to model processes as realistically as possible. When using ABM in data analysis, however, each additional rule leads to additional parameters that need to be estimated. Therefore, complex models, even if they are more externally valid, are not recommended as a starting point.

In the example study, we started with a very simplistic model with only two parameters (media and social impact). Only after the parameters for this model were estimated and the validity and reproducibility of results was established, we added rules to moderate these effects and extend the model to the six still simple rules described in this paper. Based on experiences in this modeling process and in unsuccessful alternative approaches, we advise not to aim for complex models but to get simple models to work, as is generally advised for ABM (Railsback & Grimm, 2012). When they converge and produce valid results, these simple models may gradually be extended. Other than in regression models, where the inclusion and exclusion of covariates may be tried at virtually no costs, the decision to

include additional variables, rules, and conditions have to be considered thoroughly in the development of an ABM for data analysis. Thankfully, the slow computation of these models leaves much time to ponder possible next steps.

Alternative ways of parameter estimation

The second major challenge when using ABM as an analytical tool lies in the efficient and reliable estimation of model parameters. While the general direction of effects is generally deduced from theoretical considerations and previous research, their quantification for simulations is non-trivial. In the example study, we used a relatively simple genetic algorithm to narrow down all parameters to their optimal values. When parameters are strongly interdependent (i.e.: an increase in baseline susceptibility goes together with a decrease in the effect of attitude certainty), this estimation may be time-consuming as the algorithm but slowly converges on optimal values for co-evolving parameters. In the process of narrowing down the parameters, however, a genetic algorithm offers valuable insights to the relative importance of parameters and their co-dependence. Specifically, we observed that the optimal values for parameters that are of high importance to the overall fitness of the model are found more quickly than less important parameters. In the present study, while the baseline social susceptibility converged rapidly toward its optimal value, media susceptibility remained uncertain and only began to converge when all other parameters were close to their optimum (Fig. 3).

Next to genetic algorithms, we have outlined two other possible approaches to efficient parameter estimation. The more elegant approaches use Bayesian models based on Markov chains to iteratively improve posterior assumptions on the distribution of parameters and errors to fit the data. Such an approach is used, for example, in the Siena algorithm (Snijders, 2017a) that estimates the probability functions underlying network dynamics. In order to apply such an approach to parameter estimations for an ABM, however, founded assumptions on all error variances and the distribution of parameters have to be made. Additionally, it must be possible to infer the value of an individual parameter if all other parameters and the data are known to inform the next stage in the Markov chain. Due to the chaotic non-linear progression of agent-based simulation, this condition is hardly ever met. If set up correctly, however, such an algorithm does not only provide sensible values for all parameters and their standard errors but also assesses the rate and goodness of convergence (Snijders, 2017b).

A more straightforward if less elegant approach to multiple parameter estimation is an iterative reduction of parameter intervals using a brute-force grid search. For this approach, a range of possible values for each parameter is defined, from which a small number (e.g. N=10) of equally spaced values are extracted. Running the simulation for all parameter value combinations and evaluating the explanatory power of each simulation will reveal local and global maxima. Subsequently, the ranges for each parameter may be narrowed down to values surrounding the optimum. By iteratively reducing the possible range for each parameter in this way, optimal values are eventually found with reasonable precision. This approach, however, is only feasible for simple models as the number of possible combinations and the required calculation time grows exponentially with the number of parameters. For the five parameters, used in the example study, 100'000 simulations would have to be run at each iterative stage of which at least five are required to locate the optimum down to four decimals. Each additional parameter would require tenfold computational cost. This approach is, therefore, useful for the estimation of one or two parameters in preliminary analyses but should be replaced by a more elegant solution as model complexity increases.

In addition to these three main routes toward parameter estimation, a wide range of optimization algorithms have been and are being proposed in different scientific fields to find parameter values. These include traditional approaches such as simulated annealing (van Laarhoven & Aarts, 1987), Tabu Search (Glover, 1986), or ant colony optimization (Dorigo, 1992) as well as recent applications of artificial intelligence and neural networks. When considering the application of any of these approaches to the problem outlined in this study, one has to keep in mind that looking up the value for any given parameter combination takes several seconds and is the time-limiting step in this calculation. The use of algorithms requiring the computation of large numbers of parameter sets per iterative stage is, therefore, not advised.

Since the approach of using ABM in linkage analysis does not hinge on any particular approach to parameter estimation and may well work with any algorithm able to find global maxima in an oblique parameter space, future implementation may use other and maybe faster algorithms, making the approach more efficient and easy to handle. The use of genetic algorithm in this paper was mainly owed to the simplicity of this approach both in application and explanation.

Overfitting

The third methodological challenge lies in the evident overfitting of parameters to the sample. Regardless of the approach, parameter estimation inevitably leads to a single best value for each parameter. Without reliable standard errors, these values may only be regarded as optimal for the specific sample used in the analysis and do not allow for any inductive interpretation. In addition, their applicability to other samples is unknown. In order to reduce overfitting and to determine the degree to which overfitting may pose a problem, we propose to use bootstrapping and cross-validation. By drawing several different samples of respondents from the available data, the range of true parameters may be estimated. Furthermore, by cross-validating the results to check the appropriateness of parameter sets for cases that were not included in the estimation it is possible to gauge the problem overfitting poses to the generalizability of results.

While cross-validation may be done in a single run of the model, bootstrapping is a time-consuming process. For complex parameter estimations in ABM, which may take hours to complete, repeated re-calculation from the original priors is not advised. Instead, we proposed two methods - integrated and post-hoc bootstrapping - which are considerably faster. In post-hoc bootstrapping, the optimal estimates are used as priors for the estimation in bootstrapping samples, which speeds up convergence. In integrated bootstrapping, the genetic algorithm uses different bootstrapping samples in each generation, leading to parameter estimates that do not converge in one point but in a range of optimal values.

Refined approach: Integrated sub-sampling

In order to tackle the challenge of overfitting and reduce the time-consuming parameter estimation using a genetic algorithm, we propose to use a slightly refined procedure that combines integrated bootstrapping and cross-validation. In this procedure, the sample of respondents is initially divided into two equally sized groups, which are used as training and test data for cross-validation. Since the cases are drawn without replacement, they are not exactly bootstrapping samples but just sub-samples of the data at hand. The genetic algorithm is used for only one generation on the training data before assessing the explained variance both in the training and test data. After this assessment, the complete sample is randomly split again to do the next generation of parameter estimation. Using this approach, the parameters are never fitted to one specific sample but are adapting to several possible sub-samples simultaneously.

As opposed to the procedure used in the study above, there are several advantages to this refined method. First, as the best parameter sets are not the ones adapted to one specific sample but the ones equally applicable to different sub-samples, overfitting and pre-emptive convergence are prevented. Consequently, the speed of co-evolution of interdependent parameters is enhanced. Second, the approach offers constant monitoring of overall model improvement by repeated cross-validation. Third, the reduction of the training sample by 50% leads to a reduction of computation time by 75%, as the model used in this study has an $O(n^2)$ order with regard to sample size. Fourth, the convergence of parameters may be computed from the ratio between the standard deviation of parameter means in previous generations and the standard deviation of parameters in the current generation. If this ratio drops below 1.0, the variation between generations has become lower than the variation within the current generation. This measure is similar to the convergence t-ratio proposed for stochastic simulations (Snijders, 2017b, 20).

To test this refined approach, it was applied to the example study and compared with the original procedure detailed above. Convergence was measured as the harmonic mean of ratios of past and current standard deviations for all parameters and dropped below 1.0 within 100 generations for both panel intervals. Within 600 generations, the parameter values stabilized at the values found to be optimal for the whole sample (Fig 6).
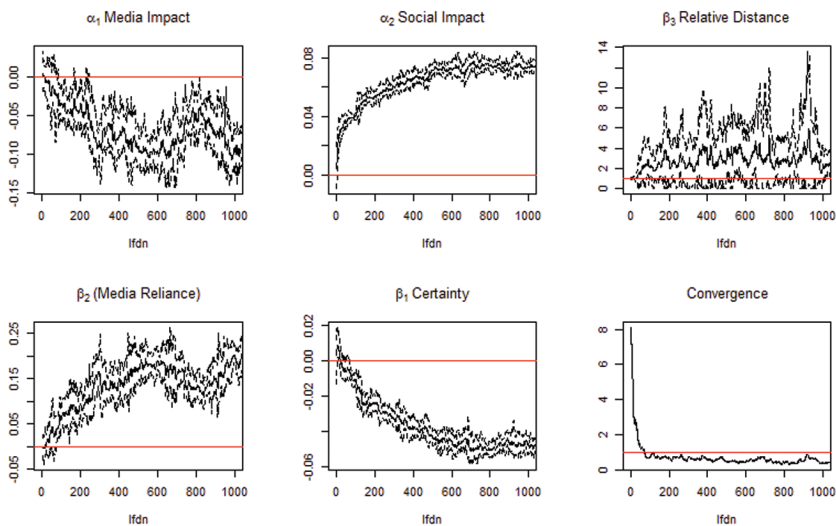

Figure 6. Parameter estimation for the interval between the first and second panel wave using integrated bootstrapping.

In comparison with the original procedure, this approach was three times as efficient in finding optimal parameters. Accounting for the fact that only a quarter of time was required for each generation of the genetic algorithm, results were obtained in less than one tenth of the time required above.

When it comes to parameter values, the results are nearly identical to the ones found by exact parameter estimation and post-hoc bootstrapping. In figure 7, the distribution of parameter values in integrated sub-sampling and post-hoc bootstrapping are compared with the optimal values for all cases for the first panel interval. While integrated sub-sampling was able to pin-point the optimal value of each parameter, it has to be noted that the standard errors of parameter estimations were likely to be underestimated in this approach. This may be due to the fact that the genetic algorithm was never able to completely fit parameters to extreme samples but was stabilized by the quick alteration of samples. Owing to this stability, we also
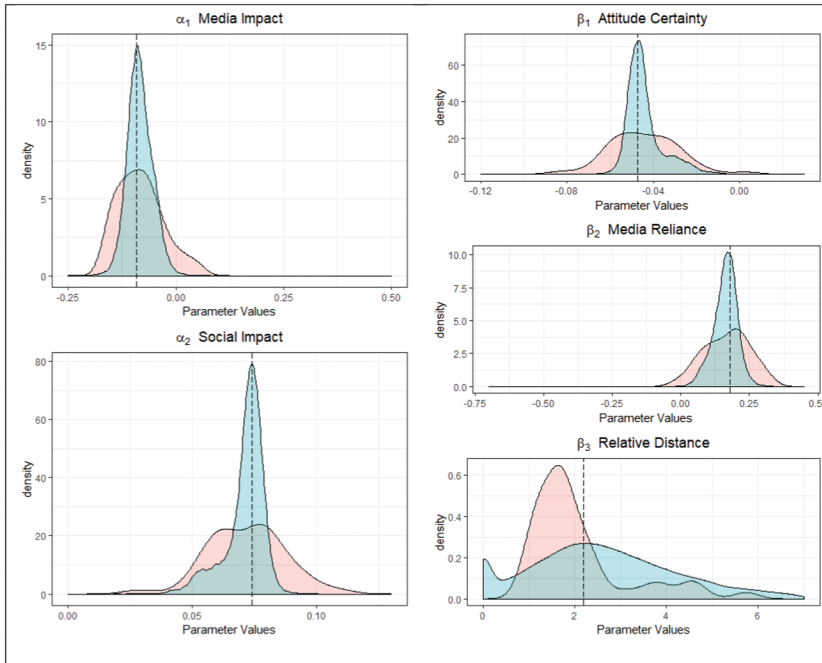

Figure 7: Comparison of parameter variance after post-hoc bootstrapping and integrated sub-sampling.

Notes: Red: Distribution of parameters after post-hoc bootstrapping (cf. Fig. 4); Blue: Distribution after integrated sub-sampling; Dashed line: Optimal value for all 1721 cases.

found a stable share of explained variance in the respective test-samples in cross-validation (M=9.62%; SD=1.12%).

Discussion

In this paper, we introduced ABM as analytical method to estimate unobserved short-term media effects in linkage analyses. Using data from panel surveys and parallel analyses of media content to define an agent-based simulation, short-termed dynamics between the panel waves may be modeled and estimated on an individual level. This approach, therefore, utilizes the unique ability of ABM to reconcile macrosocial theories with empirical data on individual level (Waldherr & Wettstein, 2019) as it allows researchers to glance into the black box of long-term opinion dynamics and to pin-point immediate media effects causing them.

Although the approach is methodologically challenging and has proven to be time-consuming even for simple models, it is a promising application of ABM in communication research. Unlike traditional linkage analyses, which use aggregated scores of media content between panel waves to assess media effects, this approach has the potential to model the immediate effect of media coverage on respondents, down to single articles.In the example study, we used a genetic algorithm to estimate all model parameters and bootstrapping to estimate their generalized probability distribution. While the computation took several days to complete, the results are striking as they reveal not only the relative importance of social and media bias in the run-up of a referendum but also showed that attitude certainty, media reliance, and ideological distance to peers are important moderators. In addition, the explanatory power of 8-10% explained variance of attitudes in the next panel wave exceeded the power of traditional linkage analyses using the same or similar datasets.

Limits and Scope
The main limitation of the proposed approach lies in the necessity of at least two panel waves. Unlike regression-based linkage analyses which may also be used to explain patterns in cross-sectional surveys by past media coverage, two measurements are required for a majority – but not all – respondents. In the example study, the inclusion of panel dropouts even increased the explanatory power as even incomplete additional agents add to the environment of complete cases. A further soft limitation lies in the time-consuming process of parameter estimation.

Since the computation time increases with each additional rule and parameter, models may not be too complex. We have shown, however, that computation time may be cut down by integrated sub-sampling. Considering possible future refinements of the approach, the possibility to find more efficient optimization algorithms, and the rapid increase of computational power, model complexity and time may not be an issue in the near future.

Bearing these limitations in mind, the scope of the proposed approach is applicable to a multitude of research questions which revolve around the accumulation of short-term effects in the long run. It is, thereby, not limited to a single outcome variable but may be used to model the dynamics on multiple dependent variables in response to media content. In the example study, only the attitude of respondents was allowed to vary over time. It would, however, be feasible to also model changes in attitude certainty (i.e.: an increase in attitude certainty if social and media bias are in line with the agent's attitude) or changes in media reliance as a function of attitudinal mismatch with media content. Accordingly, any number of media effects on an individual agent may be modeled in an ABM and put to the test using empirical data. This opens a new opportunity for media effects researchers to study the accumulation of media effects that have already been established in experimental studies but are not detectable by traditional observations in the field. By virtue of the high flexibility of this approach and the possibility to define rules for moderating effects and weighting factors, even complex interactions between media users and media content may be modeled and investigated.

In conclusion, this paper introduces a novel and promising computational method for media effects research. Combining traditional linkage analysis with the flexibility and complexity of agent-based simulations, this approach allows for the estimation of short-term individual media effects in panel surveys. In an example study, the scope and explanatory power of this novel approach was shown to exceed traditional methods of linkage analysis and to allow for a differentiated analysis of media effects on individuals. While the method has proven methodologically challenging and time-consuming, it is a valuable tool that may solve the dilemma of media effects research and close the gap between experimental and survey-based findings.

Notes

1 See https://github.com/Tarlanc/ABM_PanelWaves for full code and data for all the models described in this paper

2 For the conversion of ZIP codes to positions in kilometers, the official Swiss coordinate system (https://map.geo.admin.ch/) was used which defines the capitol in Bern at 600km east and 200km north.

3 The times required for the computation of the model were obtained on a desktop computer with an intel core i7 processor.

References

Bartlett, P. L., Boucheron, S., & Lugosi, G. (2002). Model Selection and Error Estimation. *Machine Learning*, *48*(1/3), 85–113. doi: 10.1023/A:1013999503812

Bousquet, F., Cambier, C., Mullon, C., Morand, P., & Quensiere, J. (1995). Sumulating Fishermen's Society. In N. Gilbert & J. Doran (Eds.), *Simulating societies: The computer simulation of social phenomena* (2nd ed., pp. 143–164). London: UCL Press.

Charbonneau, P. (2002). An Introduction to Genetic Algorithms for Numerical Optimization. *Advance online publication*. doi: 10.5065/D608638S

Chmiel, A., Sobkowicz, P., Sienkiewicz, J., Paltoglou, G., Buckley, K., Thelwall, M., & Hołyst, J. A. (2011). Negative emotions boost user activity at BBC forum. *Physica a: Statistical Mechanics and Its Applications*, *390*(16), 2936–2944. doi: 10.1016/j.physa.2011.03.040

Dorigo, M. (1992). *Optimization, Learning and Natural Algorithms*. Milano, IT: PhD Thesis, Politecnico di Milano.

Efron, B. (1981). Nonparametric estimates of standard error: The jackknife, the bootstrap and other methods. *Biometrika*, *68*(3), 589–599. doi: 10.1093/biomet/68.3.589

Gilbert, N., & Troitzsch, K. G. (2005). *Simulation for the social scientist*. Maidenhead, UK: Open University Press.

Gilbert, N., Ahrweiler, P., Barbrook-Johnson, P., Narasimhan, K. P., & Wilkinson, H. (2018). Computational modelling of public policy: Reflections on practice. *Journal of Artificial Societies and Social Simulation*, *21*(1), 14. doi: 10.18564/jasss.3669

Glover, F. (1986). Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, *13*(5), 533–549. https://doi.org/10.1016/0305-0548(86)90048-1

Hayes, A. F., & Scharkow, M. (2013). The relative trustworthiness of inferential tests of the indirect effect in statistical mediation analysis: Does method really matter? *Psychological Science*, *24*(10), 1918–1927. doi: 10.1177/0956797613480187

Kepplinger, H. M., Brosius, H.-B., & Staab, J. F. (1991). Opinion Formation in Mediated Conflicts and Crises: A Theory of Cognitive-Affective Media Effects. *International Journal of Public Opinion Research*, *3*(2), 132–156. doi: 10.1093/ijpor/3.2.132

Koskinen, J. H. & Snijders, A. B. (2007). Bayesian inference for dynamic social network data. *Journal of Statistical Planning and Interference 13*, 3930-3938.

Latané, B. (1981). The Psychology of Social Impact. *American Psychologist*, *36*(4), 343–356.

Macy, M. W., & Willer, R. (2002). From Factors to Actors: Computational Sociology and Agent-Based Modeling. *Annual Review of Sociology*, *28*(1), 143–166. doi: 10.1146/annurev.soc.28.110601.141117

Matthes, J. (2012). Exposure to Counterattitudinal News Coverage and the Timing of Voting Decisions. *Communication Research*, *39*(2), 147–169. doi: 10.1177/0093650211402322

Miller, A. H., Goldenberg, E. N., & Erbring, L. (1979). Type-Set Politics: Impact of Newspapers on Public Confidence. *American Political Science Review*, *73*(01), 67–84. doi: 10.2307/1954731

Muis, J. (2010). Simulating political stability and change in the Netherlands (1998-2002): An agent-based model of party competition with media effects empirically tested. *Journal of Artificial Societies and Social Simulation*, *13*(13), 4. doi:10.18564/jasss.1482

Muis, J., & Scholte, M. (2013). How to find the 'winning formula'? Conducting simulation experiments to grasp the tactical moves and fortunes of populist radical right parties. *Acta Politica*, *48*(1), 22-46. doi:10.1057/ap.2012.21

Noelle-Neumann, E. (1974). The Spiral of Silence: A Theory of Public Opinion. *Journal of Communication*, *24*(2), 43–51.

Nowak, A., & Latané, B. (1995). Simulating the emergence of social order from individual behavior. In N. Gilbert & J. Doran (Eds.), *Simulating societies: The computer simulation of social phenomena* (2nd ed., pp. 63–84). London: UCL Press.

Nowak, A., Szamrej, J., & Latané, B. (1990). Prom Private Attitude to Public Opinion: A Dynamic Theory of Social Impact. *Psychological Review*, *97*(3), 362–376.

Peter, J. (2003). Country Characteristics as Contingent Conditions of Agenda Setting. *Communication Research*, *30*(6), 683–712. doi: 10.1177/0093650203257844

Railsback, S. F., & Grimm, V. (2012). *Agent-based and individual-based modeling: A practical introduction*. Princeton, NJ: Princeton University Press.

Scharkow, M., & Bachl, M. (2016). How Measurement Error in Content Analysis and Self-Reported Media Use Leads to Minimal Media Effect Findings in Linkage Analyses: A Simulation Study. *Political Communication*, *9*(4), 1–21. doi: 10.1080/10584609.2016.1235640

Schemer, C. (2014). Media Effects on Racial Attitudes: Evidence from a Three-Wave Panel Survey in a Political Campaign. *International Journal of Public Opinion Research*, *26*(4), 531–542. doi: 10.1093/ijpor/edt041

Schemer, C., Wirth, W., & Matthes, J. (2012). Value Resonance and Value Framing Effects on Voting Intentions in Direct-Democratic Campaigns. *American Behavioral Scientist*, *56*(3), 334–352. doi: 10.1177/0002764211426329

Scheufele, B. (2008). The dilemma of media effects research. A logic for modeling media effects on meso- and macro-level units both in theoretical und methodical respect [German]. *Publizistik*, *53*(3), 339–361. doi: 10.1007/PL00022227

Schuck, A. R. T., Boomgaarden, H. G., & Vreese, C. H. de. (2013). Cynics All Around? The Impact of Election News on Political Cynicism in Comparative Perspective. *Journal of Communication*, *63*(2), 287–311. doi: 10.1111/jcom.12023

Snijders, T. A. B. (2017a). Stochastic Actor-Oriented Models for Network Dynamics. *Annual Review of Statistics and Its Application*, *4*(1), 343–363. doi: 10.1146/annurev-statistics-060116-054035

Snijders, T. A. B. (2017b). Siena Algorithms. Retrieved from https://www.stats.ox.ac.uk/~snijders/siena/Siena_algorithms.pdf

Van Laarhoven, P. J. M., & Aarts, E. H. L. (1987). Simulated annealing. In P. J. M. van Laarhoven & E. H. L. Aarts (Eds.), *Mathematics and its applications: Vol. 37. Simulated annealing: Theory and applications* (1st ed., pp. 7–15). Dordrecht: Springer. doi: 10.1007/978-94-015-7744-1_2

Voinov, A., Kolagani, N., McCall, M. K., Glynn, P. D., Kragt, M. E., Ostermann, F. O., Ramu, P. (2016). Modelling with stakeholders – Next generation. *Environmental Modelling & Software*, 77, 196–220. doi: 10.1016/j.envsoft.2015.11.016

Waldherr, A. (2014). Emergence of news waves: A social simulation approach. *Journal of Communication*, *64*(5), 852-873. doi:10.1111/jcom.12117

Waldherr, A. & Wettstein, M. (2019). Bridging the Gaps. Using Agent-Based Modeling to Reconcile Data and Theory in Computational Communication Science. *International Journal of Communication* 13, 3976-3999.

Wettstein, M. (2012). Frame Adoption in Referendum Campaigns: The Effect of News Coverage on the Public Salience of Issue Interpretations. *American Behavioral Scientist*, *56*(3), 318–333. doi: 10.1177/0002764211426328

Wirth, W., Matthes, J., Schemer, C., Wettstein, M., Friemel, T., Hänggli, R., & Siegert, G. (2010). Agenda Building and Setting in a Referendum Campaign: Investigating the Flow of Arguments Among Campaigners, the Media, and the Public. *Journalism & Mass Communication Quarterly*, *87*(2), 328–345.

About the author

Martin Wettstein (Ph.D.), is a post-doctoral researcher in the Department of Media Psychology and Effects at the University of Zurich. In research and teaching, he focuses on quantitative methods, public opinion dynamics, and political communication.

Correspondence address: m.wettstein@ikmz.uzh.ch