24. Agent-based computational models¹

Andreas Flache and Carlos A. de Matos Fernandes²

1. INTRODUCTION

Agent-based computational modeling (ABCM) plays a central role in analytical sociology (Hedström and Bearman, 2009). Analytical sociology (AS) aims to explain social phenomena with precisely described and empirically realistic social mechanisms. Mechanistic explanations in AS target, in particular, complex micro-macro interactions in which individual behavior at the micro level of society can, in interaction with meso- and macro-level structures, bring about macro outcomes individuals neither necessarily anticipate nor desire. Examples are residential segregation (Clark and Fossett, 2008), opinion polarization (Flache, Mäs, et al., 2017), or dysfunctional status hierarchies (Mark et al., 2009). AS aspires to uncover why this occurs in real-life social settings by modeling the underlying social mechanism at a level of precision that allows logical deduction and empirical refutation of both model assumptions and predictions. With this ambition, AS is not fundamentally different from, for example, empirically oriented Rational Choice Sociology. However, despite the considerable flexibility of modern conceptualizations of Rational Choice Sociology (Wittek et al., 2013), AS is different in that it does not commit to a particular set of ingredients for modeling individual behavior (Manzo, 2014, pp. 21–27). ABCM's capacity to combine high theoretical flexibility with analytical precision is exactly what makes it so attractive to analytical sociologists.

Computer simulation is the main approach in ABCM for analyzing whether, and if so, how, and under which conditions exactly the social mechanism formulated in a computational model can generate the phenomenon the researcher wants to understand (Epstein, 2006). While ABCM is not the only method for formalizing social mechanisms, it is particularly compatible with AS (Macy and Flache, 2009; Manzo, 2014, 2021, *forthcoming*) because it combines formal precision and the ability to model and analyze complex micro-macro interactions (Mäs, Chapter 4, this volume) and theoretical flexibility in a unique way.

ABCM forces a modeler to make the explanatory mechanism fully explicit. The need to translate a conceptual mechanistic model into an executable computer program requires describing algorithmically how agents arrive at changes in their actions, cognitions or emotions

¹ This is an open-access chapter distributed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 Unported (<https://creativecommons.org/licenses/by-nc-nd/4.0/>). Users can redistribute the work for non-commercial purposes, as long as it is passed along unchanged and in whole, as detailed in the License. Edward Elgar Publishing Ltd. must be clearly credited as the rights holder for publication of the original work. Any translation or adaptation of the original content requires the written authorization of Edward Elgar Publishing Ltd.

² The first author gratefully acknowledges financial support by the Netherlands Organization for Scientific Research (NWO) under the 2018 ORA grant ToRealSim (464.18.112). Both authors acknowledge that this study is part of the research program Sustainable Cooperation—Roadmaps to Resilient Societies (SCOOP) funded by NWO and the Dutch Ministry of Education, Culture and Science (OCW) in its 2017 Gravitation Program (grant number 024.003.025). Further, we are grateful to Gianluca Manzo for thoughtful feedback and editorial support, as well as to the members of the Norms and Networks Cluster (NNC) at the University of Groningen, Department of Sociology, for an inspiring and helpful discussion of an earlier version of this chapter.

interdependently with their local and global social context. This is different from the approach taken by game theory, where individual actions are seen as part of a strategy profile in a game that satisfies axiomatic requirements for strategically rational choice, while the exact cognitive process that leads actors to make those choices is left implicit (Breen, 2009). Formal approaches with less demanding cognitive assumptions like evolutionary game theory (Bendor and Swistak, 2001), Markov chain models (Young, 1998) or socio-physics (Castellano et al., 2009) are similarly capable of addressing complex micro-macro interactions and offer in addition powerful analytical tools that can yield more general insights into model implications than ABCM. However, these approaches may require a modeler to press their substantive theory of a social mechanism into the Procrustean bed of a particular formal framework. From the AS perspective this is undesirable because realistic models of social mechanisms should only be constrained by the theoretical and empirical information at hand, as well as the requirement of logical consistency and precision. Keuschnigg et al. (2018, p. 6) express this view as follows:

It is not enough that the theory or model predicts the phenomena to be explained; an explanation must represent the essential features of the actual process that produced the phenomena to be explained. [...] from the AS perspective, considerations of elegance, simplicity, or tractability never should override the aim of accurately describing the mechanisms that actually produced the phenomena to be explained.

Analytic precision, ability to capture complex micro-macro interactions, and flexibility to accommodate empirically realistic assumptions suggest for many adherents of AS that ABCM is a highly suitable approach to provide mechanistic explanations. But this does not mean that anything goes. Other than analytical approaches like game theory or socio physics, ABCM typically does not come with the possibility to mathematically prove that model results follow correctly from model assumptions. Empirically plausible models that mimic real phenomena are no guarantee that the process by which the model generated the result has been sufficiently understood by the modeler to exclude artefacts created by inessential model details or perhaps programming errors. ABCM is most needed when the implications of a seemingly simple mechanism are too complex for a researcher to be sufficiently understood without formal modeling. But this also implies that a model built to capture a complex social mechanism can be a complex system in itself that needs to be understood well in the first place. Without the full comprehension of the dynamic processes generated by an ABCM, little has been gained if these dynamic processes resemble the social phenomena the researcher wants to explain. In other words, explaining the outcomes generated by the ABCM itself is different from but logically precedes explaining a social phenomenon through that ABCM.³ Yet, there is still a major gain from creating a complex ABCM for explaining a complex phenomenon. Other than social reality, the computational model is fully transparent and fully under the control of the researcher, allowing for systematic analysis and experimentation with it. Thus, the potential of ABCM to aid AS can be fully realized if combined with adherence to a set of good practices for systematic deductive model development and analysis.

This chapter aims to give readers an easily accessible introduction to ABCM in the context of AS. Gaining practical experience is crucial for understanding the benefits and potential

³ We thank Gianluca Manzo for having inspired this formulation.

pitfalls of ABCM. A core element of the chapter is, therefore, the development of a modeling example the reader can replicate themselves with *NetLogo* (Wilensky, 1999), a freely available software. We start from Schelling's famous model of segregation, the same "toy-model" (Reutlinger et al., 2017) that Hedström and Bearman (2009) used to illustrate how ABCM helps to understand complex micro-macro interactions.⁴ We extend this model toward including more empirically informed assumptions. This demonstrates the use of the good practices we propose for ABCM, showing how it can be challenging in ABCM to understand what the consequences of adding realism to a model are, and how navigating between simpler, less realistic models, and more complex models with a higher degree of realism is needed for full comprehension of model outcomes.

In what follows, we first describe the main ingredients of an ABCM (section 2). Section 3 presents the modeling example and good practices. We conclude the paper with a broader reflection of the role of ABCM in AS.

2. A BIRD'S EYE VIEW ON ABCM

What follows is a brief look at the main components of an ABCM. Comprehensive overviews and discussions of ABCM in sociology, and AS in particular, can be found in the literature (Macy and Flache, 2009; Squazzoni, 2012; Manzo, 2014; Bianchi and Squazzoni, 2015).

The core element of an ABCM are **agents**, software objects representing individual social actors construed simultaneously as *autonomous*, *interdependent*, *heterogeneous* and *embedded* (Wooldridge and Jennings, 1995). Agents comprise state variables, representing fixed characteristics like their sex or skin color, but also beliefs, opinions, emotions, or actions that agents change autonomously, depending on their current state and on external inputs. From the perspective of AS, agent-level action variables must be construed such that the actions of all agents jointly constitute **macro-level outcomes**, mirroring the empirical social phenomenon of interest. An example is residential segregation, a macro-outcome computed from the distribution of individual agents over the space of an artificial city, which in turn is the result of agents' residential choices and moving actions (Benenson et al., 2009). The initial value of state variables, representing internal agent states, beliefs, opinions or initial residential locations, as well as initial macro-level characteristics like the size of an agent population, can be seen as **modeling the social situational context** of which the modeler wants to understand why and how it is related to the social outcome of interest.

A **model of agent behavior** in an ABCM comprises both assumptions about agents' internal states and actions as well as assumptions about their social interdependence. In modeling decision-making, ABC modelers typically impose comparatively low demands on the cognitive ability of agents, but explicate agent's decision-making process, motivated by critique of the descriptive inaccuracies of the axiomatic behavioral model of perfectly rational

⁴ A slightly updated version of the model Hedström and Bearman (2009) used is available in the NetLogo model library under the name *segregation.nlogo* (Wilensky, 1997). The models presented in this chapter are considerably further developed to include empirically more realistic assumptions. Following an important good practice in ABCM, we make our code available for replication. We provide two *NetLogo* models of segregation, called *SegregationExtended* and *SegregationDiscreteChoice*, respectively, available for download at <https://www.comses.net/codebase-release/db473a1d-c803-4e9d-8b2f-d9389784a063/>. To use the models, first download and install the software NetLogo (<https://ccl.northwestern.edu/netlogo/download.shtml>, version 6.1.1.).

choice (Macy and Flache, 2009). Bounded rationality (Simon, 1982), heuristic decision-making (Goldstein, 2009), simple reinforcement learning (Macy and Flache, 2002) or more sophisticated forms of learning (Zschache, 2018) are common approaches used in modeling individual-level behavior in an ABCM.

While autonomous agents control their own internal states, their autonomy is constrained by **interdependence**. Each agent's actions and possibilities for future actions depend on what happens in the agent's local environment, and these actions in turn have consequences that alter their own environment and that of others. Important for many applications in the social sciences, agents are locally interdependent, for example in a social network that links them with a (typically small) subset of the overall population, such that macro-level outcomes are an emergent property of local interactions. Local interaction is often seen as a key ingredient to a realistic description of social dynamics in large populations (Axelrod, 1997). In addition, ABCMs can include assumptions about how agents change their position in a geographical space or break off relations and establish new ones with other agents (Centola et al., 2007; Bianchi et al., 2020).

ABCMs have been applied to study the emergence and dynamics of a wide range of social phenomena, including residential segregation (Schelling, 1971; Clark and Fossett, 2008; Hegselmann, 2017), cultural differentiation (Axelrod, 1997; Flache and Macy, 2011), norm diffusion (Mäs and Opp, 2016), opinion polarization (Baldassarri and Bearman, 2007; Flache et al., 2017) or the diffusion of innovations (Manzo et al., 2018). In the next section, we use a model of residential segregation to illustrate good practices and potential pitfalls generally applicable to ABCM.

3. AN ABCM OF UNINTENDED RESIDENTIAL SEGREGATION: FROM A TOY-MODEL TO A (SOMEWHAT MORE) REALISTIC MODEL

Schelling (1971) famously proposed an extremely simple but highly illustrative social mechanism to understand how strong ethnic segregation could arise in a world where individuals do not actually want it. Schelling's simple computational model will be taken as the starting point for our modeling example, but readers should be aware that independently different models implementing essentially the same mechanism have been proposed (Sakoda, 1971; Hegselmann, 2017). The core part of this mechanism is a *relocation cascade*. Even if individuals are content with living in a mixed neighborhood, they also want at least a certain minimal fraction of co-ethnics nearby. As soon as some individuals feel outnumbered, they leave and settle elsewhere. In the process, movers unintentionally make it more likely that other members of their group follow suit, because after they leave, their group becomes even more outnumbered in their prior neighborhood. Multiple modeling studies have shown how this simple micro-level process can induce a self-reinforcing cascade in which more and more relocations occur until neighborhoods end up being highly segregated (for an overview, see Hegselmann, 2017).

3.1 Social Context and Macro-Outcomes: Which Social Phenomenon Is to Be Explained?

AS wants to provide precise and realistic mechanistic explanations of concrete and puzzling empirical macro-phenomena. Schelling's model, in particular, addressed what he saw as a

puzzle. Ethnically diverse cities in the US (and elsewhere) are highly ethnically segregated, although empirical studies suggest that most people would prefer living in mixed neighborhoods (Clark and Fossett, 2008).

ABCM forces a researcher to formulate precise, formal descriptions of all elements of a model in computer language. However, just expressing a social mechanism in computer code is not enough.

Good practice #1: Formulate precise descriptions of the social phenomenon to be generated in terms of the macro-level outcome, its situational context and of the postulated generative mechanism before writing computer code.

Good practice #1 helps model builders keep their focus on the research problem and the theoretical assumptions that constitute a postulated mechanism rather than on technical details of their computational implementation.

Example of a research goal in ABCM:

Show that a relocation cascade can generate higher levels of segregation than needed to satisfy individuals' ethnic preferences. Further, identify the conditions under which this outcome obtains as robust regularity.

The macro-level outcome of interest is here specified as segregation, the proposed explanatory mechanism is that of a relocation cascade and a key element of the situational context is that ethnic preferences are mild in the sense that individuals would be content with less segregation than they get. The *NetLogo* model *SegregationExtended* gives specific technical elaborations for each of these core substantive concepts.

3.2 A Slightly More Realistic Toy-Model

SegregationExtended encompasses Schelling's classic toy-model as a special case.⁵ We start here from a version with slightly more realistic assumptions, assuming in particular larger neighborhoods and a population consisting of four groups roughly resembling the ethnic composition of a contemporary large US city. Agents in the program are called *households*. The **model of agent behavior** is extremely simple. Households' only "action" is their choice of a residential location, represented by which cell of a rectangular checkerboard-like "world" (hereafter: cellular world) the household chooses to locate to at a given point of time. Further, households have the state variable *ethnicity*, assigned as a fixed characteristic upon initialization of the simulation.

⁵ This chapter cannot give an introduction to *NetLogo* itself. Further explanation of the model can be found under the tabs "Code" and "Info", respectively, on the model's interface and in many explanatory comments in the code added by the authors (starting with ";;"). The code of the models *SegregationExtended* and *SegregationDiscreteChoice* has not been optimized for efficiency, but instead for transparency for users new to *NetLogo*. A tutorial on the *NetLogo* website (<https://ccl.northwestern.edu/netlogo/docs/>) provides an accessible introduction to *NetLogo* itself. Further comprehensive introductions in the literature are Wilensky and Rand (2015) and Railsback and Grimm (2019).

By default, *SegregationExtended* assumes a **situational context** with four different groups ($ethnicity \in \{1,2,3,4\}$). Following Clark and Fossett (2008), the agent population resembles in this setting roughly the composition of a typical large US city with a white majority, and African Americans, Asians, and Hispanics with proportions of 60%, 20%, 10%, and 10%, respectively. In the model, these groups are called Red, Green, Blue, and Yellow, respectively. Switching *fourGroups* “Off” creates two equally large groups Red and Green ($\in \{1,2\}$), as often used in the literature (Hedström and Bearman, 2009). In both settings, households are in the initial condition randomly distributed over the cellular world. While this may be empirically implausible given the high levels of ethnic segregation observed in many multi-ethnic societies, this initial condition is of high theoretical interest for assessing whether and under which conditions relocation cascades give rise to high levels of segregation, especially if households do not prefer nor live in segregated neighborhoods initially.

Figure 24.1 shows examples of states of the cellular world with four groups as they are shown on the *NetLogo* interface of *SegregationExtended*. The probability that a cell is occupied by a household rather than being empty is a further important element of the situational context in this model, defined by the program slider *density*.

Households’ ethnic preferences are represented by the threshold T for the minimal proportion of co-ethnics a household wants to have (slider *%-similar-wanted*). In *SegregationExtended*, T is measured in percentage points and equal for all households, represented by a so-called “global” variable known to all households in the program. The model makes exactly one simple assumption about households’ cognition (see procedure *update-households*). If the proportion of co-ethnic households in the neighborhood falls below T , the household is unhappy. The model of action is equally simple (see procedure *move-unhappy-households*). If a household is unhappy, it searches randomly a vacant spot, starting from its current location, and stays in it as soon it has found one.

Figure 24.2C shows how households are **locally interdependent** with their neighbors. Typically, Schelling’s (1971) model is presented with the assumption of very small neighborhoods with up to eight local neighbors (see also Hedström and Bearman, 2009), but this may seem empirically implausible. Benenson et al., (2009), for example, propose that neighborhoods encompass the houses around a residential location that are visually noticeable. Depending on the exact spatial structure, this could be considerably more than eight households. Thus, the model assumes by default a larger neighborhood size with up to 28 neighbors (see Figure 24.2C). Yet, model assumptions about neighborhood size may vary, depending on theoretical and empirical considerations. The slider *radius* in *SegregationExtended* allows changing neighborhoods as demonstrated in Figure 24.2C, where *radius* = 3 is the default. For households located on the border of the cellular world, their neighborhood does not stretch beyond the border.⁶

The button *Go* on the *NetLogo* interface starts the procedure *go*, which controls the simulation execution. In every round of the simulation (*tick* in *NetLogo*), the procedure *move-unhappy-households* calls on all households in a random sequence and activates the procedure *find-new-spot* for unhappy households. Further, the model contains a very simple **stopping rule**. When all agents are happy the simulation ends, because an equilibrium state has been reached in which no further relocation can occur.

⁶ Often it is assumed that cellular worlds are “wrapped around” on a torus (Hegselmann and Flache, 1998) to avoid border-effects. Our world here has borders, just as real cities have.

Finally, the model reports **outcome measures** and graphical output (e.g. Figure 24.2A) to show how **macro-level characteristics** change when the micro-level processes of household relocation dynamics are at work. *Percent Similar*, the average proportion of co-ethnics among a household's neighbors, is computed for the entire population and separately for each of the four groups. Comparison of its value over time with the initial level shows segregation, indicating whether households sort into clusters that are more ethnically homogenous than in the initial randomly mixed population. Also, a comparison of its value with the threshold T tells whether households end up more segregated than needed to satisfy their ethnic preference. Further outcome measures will be introduced next.

3.3 Model Analysis: Dissecting a Toy-Model

Hedström and Bearman (2009, p. 13) demonstrated how the simplest version of Schelling's model can generate a complex micro-macro interaction. Using two equally sized groups and a small neighborhood, they showed how an increase of the ethnic preference T by only one percentage point (from 25% to 26%) remarkably changed the macro-level outcome from near random mixing to clear segregation. *SegregationExtended* slightly differs in some model details from their version, for example in the exact shape of neighborhoods or the exact number of households and density.⁷ Yet, the reader can easily check that *SegregationExtended* produces almost the same result by running simulations for a situational context similar to theirs (*radius* = 1, *fourGroups* = *Off*, 51x51 world size, 90% density). In other words, the mechanism of relocation cascades meets our research goal independently from some details of how it is programmed. This demonstrates an important good practice in ABCM.

Good practice #2: Always test whether substantively important model results are robust against small changes of theoretically irrelevant model details.

But will we find a similarly sharp transition in macro-level outcomes also for our more realistic default scenario with four groups and larger neighborhoods? Figure 24.1 shows the outcomes of a conservative test. We compared model results for $T = 25$ and $T = 35$. The figure shows two representative outcomes that emerged after 1000 simulation rounds with the default settings of *SegregationExtended*.

As it turns out, the more realistic model cannot generate a similarly sharp qualitative transition in macro outcomes. Figure 24.1 shows that an increase of ethnic preference T by 10% yields an increase of the level of segregation by about 4, a change that is very modest compared to the strong effect of a 1% increase of T when we assume two groups, *radius* = 1 and otherwise the same conditions ($T = 25$ and $T = 26$ yielded 58% and 78.7%, respectively). Further simulations with our more realistic scenario show that for $T = 25$ vs. $T = 26$, the difference in segregation is now almost negligible, with *Percent Similar* of 52.8 vs. 53.1 in two representative runs.

⁷ The situational conditions in their simulation differ from the standard settings in *SegregationExtended* as follows. They use a cellular world without borders, they assume exactly 1 250 households per group and they use a so-called Moore neighborhood with 8 neighbors (NetLogo dictionary: *neighbors*), while our *radius* = 1 imposes four neighbors (Von Neumann neighborhood). Nevertheless, outcomes are very similar. When we increased T from 25 to 26, *Percent Similar* soared from 55.6 to 77.3 in two representative runs with *SegregationExtended*.

Figure 24.1 Final states of two representative runs for different values of ethnic preference threshold (T). The proportion of Red, Green, Blue, and Yellow households is 60%, 20%, 10% and 10%, respectively. Groups are based on Clark and Fossett (2008): African American = Green; Asian = Blue; Hispanic = Yellow; White = Red. World size = 51×51 , 90% density (≈ 2345 households and ≈ 255 empty cells)

This discrepancy between the original toy-model and our slightly more realistic version serves to illustrate several important good practices in ABCM.

Good practice #3. Show the statistical distribution of model outcomes.

Good practice #4. Always formulate a candidate explanation and tentative conjecture about the expected model outcomes before running a simulation experiment.

Good practice #5. Design and conduct simulation experiments for testing the candidate explanation.

Good practice #3 is straightforward and common in ABCM research. Models typically generate stochastic processes and single runs may say little about the distribution of results. To assure robustness, we used NetLogo's *BehaviorSpace* feature⁸ to run 10 independent repetitions per condition. Results supported robustness, yielding averages and standard deviations for *Percent Similar* in the final state of 53.27 (0.63), 53.04 (1.21) and 56.32 (2.89) for $T = 25$,

⁸ The implementations of all simulation experiments shown in this chapter are embedded in the models *SegregationExtended* and *SegregationDiscreteChoice* and can be found under Tools > BehaviorSpace on the NetLogo interface.

26 and 35, respectively. In all examples shown in this chapter, there was very little variation between the outcome measures of different independent repetitions, which made further statistical tests or visualization of distributions unnecessary. This is not always the case and sometimes many more repetitions may be required.

Good practices #4 and #5 are less straightforward. However impressively an ABCM demonstrates a complex micro-macro interaction, modelers must be able to explain why. “The computer said so” is not good enough as an answer. Beautiful results from simulation studies may turn out to be hard to communicate to other researchers or practitioners and, what is more, there is no a priori guarantee that simulation outcomes are systematically generated by the theoretical mechanism the modeler wanted to implement rather than by inessential technical details or—even worse—programming errors. Every experienced ABCM practitioner knows stories of how beautifully counterintuitive insights from simulation models fell apart after bugs in the code or conceptual artefacts in the model were discovered (see, for example, a discussion in the literature about an extension of Schelling’s model: Bruch and Mare, 2006; van de Rijt et al., 2009).

It is difficult to pin down what exactly explaining model outcomes means. Ideally, mathematical proofs can be given that a simulation algorithm or a simpler formalization of essentially the same mechanism must generate the observed outcomes. Sometimes such proofs can be found or are available in the literature, e.g. for the simplest versions of Schelling’s model (Young, 1998). If they are, they form an important source of guidance for model understanding. However, often such results are simply not available (yet) or only apply for overly strong simplifications of a model. Even then, researchers should be able to produce clear causal arguments for why a model generates a certain outcome under certain conditions. Similar to theory-driven empirical research, the ultimate test is to derive expectations about model behavior in yet unexplored conditions and test these expectations against the computational evidence the model generates. Unlike empirical research, ABCM allows modelers to experiment with and fully observe every detail of the complex system the researcher created. This allows designing and running experiments which advance understanding of the model by putting tentative explanations to a critical test.

Let us apply good practices #4 and #5 to our example. Why at all does segregation increase if we raise T , and why, more in particular, is the micro-macro interaction of our more realistic model so much less dramatic than for the simpler version with two equally sized groups and $radius = 1$? Answering the first question provides a candidate explanation that helps to answer the second. The larger the proportion of initially unhappy households, the more households will move and the bigger the relocation cascades eventually will become. Further, the bigger the relocation cascades are, the more segregation arises in the final state, because every relocation tends to increase neighborhoods’ ethnic homogeneity. This yields a testable expectation. The larger the proportion of initially unhappy households, the higher segregation will eventually become. As the next experiment will show, this reveals why the micro-macro interaction is qualitatively different to the original toy-model. The crucial reason is that neighborhoods are much larger with $radius = 3$.

Figures 24.2A and B show results of a simulation experiment (Experiment 1 and 2 in *BehaviorSpace*, respectively) designed to systematically test the candidate explanation. Ethnic preference T was varied starting from $T = 0\%$ to $T = 70\%$ in steps of 1%, and the experiments were conducted for two different neighborhood sizes ($radius = 1$ vs. 3). Simulations were stopped after maximally 1000 rounds if they did not reach equilibrium before. The chart reports two outcome measures, the proportion of households who were unhappy in the initial

Figure 24.2 Effects of ethnic preferences (T) on percent similar after 1000 simulation rounds (A) and percent unhappy initially (B), averaged across 10 repetitions per level of T . Neighborhoods with radius 1 and 3 are visualized in panel C. Panels A and B are generated with Experiment 1 and 2, respectively. World-size 51×51 , 90% density (about 2345 households)

state (Experiment 2) and *Percent Similar* in the final state (Experiment 1). The addition of the initial proportion of unhappy households illustrates that in combination with systematic experimentation, properly designed outcome measures collecting model-generated “data” about the internal processes of a model are an important source of information for testing tentative explanations of model outcomes.

Results clearly support our expectations as long as T is not too high. Especially the results for $radius = 1$ show that the levels of T at which the initial proportion of unhappy agents rises sharply, coincide exactly with similarly sharp increases of *Percent Similar*. For $radius = 3$ we see instead a much smoother increase, alongside a similarly smooth increase of the initial proportion of unhappy households.

Experiments 1 and 2 partially support our candidate explanation. But why are outcomes so different for $radius = 1$ and $radius = 3$? The difference in neighborhood structures (Figure 24.2C) shows why. Households with four neighbors ($radius = 1$) want at least one co-ethnic if $0 < T \leq 25\%$, they require at least two co-ethnics if $25\% < T \leq 50\%$, three co-ethnics are needed when $50\% < T \leq 75\%$ and only four co-ethnics can satisfy a household when $T > 75\%$. This is why under $radius = 1$ the proportion of initially unhappy households goes upward in sudden jumps when ethnic preferences T change, inducing corresponding sudden jumps in segregation.⁹ However illustrative, the strong micro-macro interaction demonstrated for $radius = 1$ turns out to be an artefact caused by extremely small neighborhoods. With the larger neighborhood of $radius = 3$, a household’s initial neighbors resemble a representative sample of the overall population. Here, increasing

⁹ Figure 24.2A also shows smaller jumps for $radius = 1$ at $T = 33$ and $T = 66$. The reason is that households at the border have three rather than four neighbors.

T changes initial happiness rates much more smoothly, proportionally to how likely an average household has $T\%$ or more co-ethnic neighbors in a random sample ($N = 28$) of the overall population.

Putting good practices #4 and #5 into action, Experiments 1 and 2 tested our candidate explanation. However, this test yielded only partial confirmation. What was not anticipated was that segregation in Figure 24.2A suddenly dropped beyond a critical level of T , although initial unhappiness further increased (Figure 24.2B). For $radius = 1$, the drop occurs at about $T = 50\%$, for $radius = 3$ the critical point is at about $T = 60\%$. This brings us to the next good practice.

Good practice #6. Inspect single runs of the conditions where model behavior cannot yet be explained.

Observing the dynamics of single runs at various levels of $T > 50\%$ quickly shows what happened. It is nearly impossible that all households find vacant locations with this many co-ethnic neighbors just by a random walk. Thus, even after 1 000 simulation rounds many households remain unhappy and keep wandering around at random, strongly reducing segregation. Majority households can still form clusters to some extent, but minority households fail to do so. This explains why *Percent Similar* does not fully drop to the level of random mixing. Our candidate explanation needs refinement. If T exceeds a critical level, the proportion of unhappy households becomes too high so that the formation of stable ethnic clusters becomes impossible and segregation drops.

The example illustrates two lessons of more general importance. First, it shows how important it is to make sure one has understood unexpected model results before further complications are added to a model.

Good practice #7. Always develop an explanation of unexpected simulation results and test it, before developing the model further.

The example also demonstrates how important it is to think carefully about the stopping rule applied in a simulation experiment and to explore model behavior with single runs a priori. Had we not fixed the maximum number of rounds at 1 000, Experiment 1 would never have finished within a reasonable amount of time. That is, the other stopping condition—that all agents are happy—was practically impossible to ever be met for high levels of T . Inspecting single runs before actually experimenting made it possible to find this out in advance.

Good practice #8. Inspect single runs by sampling the whole parameter space of an experiment before designing the experiment and stopping rule.

On the whole, the somewhat more realistic model we have presented here meets the research goal formulated. We identified conditions for which it generates levels of segregation which clearly exceed those of random mixing and those needed to satisfy agents' preferences T . At the same time, our model differs in even more ways from Schelling's original toy-model than could be discussed here. This includes, in particular, the counterintuitive result that some groups actually cluster less when their ethnic preferences become more demanding. The explanation for this also makes clear why in Figure 24.2A we observe some decline of *Percent Similar* between about $T = 25$ and $T = 30$, another result not anticipated by our a priori

candidate explanation. The reader can find elsewhere how this intriguing complexity can be discovered and unraveled with further simulation experiments.¹⁰

3.4 Toward a More Realistic Model of Relocation Decision-Making

3.4.1 Model extension

In this section, a final model extension is presented in which we changed three aspects of the model of **agent behavior**, increasing the empirical plausibility of assumptions about relocation decision-making. In a real ABCM project these aspects should be introduced and analyzed step by step. For lack of space, we focus here on assessing whether the resulting model meets our research goal.¹¹

First, ethnic preferences are made heterogeneous in that households differ within and between groups in how tolerant they are toward outgroups in their neighborhood (Xie and Zhou, 2012). Second, we drop the assumption of thresholds, assuming instead that neighborhood attractiveness differs depending on how closely households' ideal neighborhood composition is matched (Zhang, 2004). Finally, the more attractive a vacant location is for a household, the more likely the household will relocate to it (Bruch and Mare, 2006; van de Rijdt et al., 2009). This last extension, in particular, connects ABCM with empirical research (Bruch and Atwell, 2015), using discrete choice models (Bruch and Mare, 2009) for which the parameters can be estimated from empirical data (van Gent et al., 2019), including experiments (Holm et al., 2016).

Ethnic preferences in *SegregationDiscreteChoice* are loosely based on empirical data. Table 24.1 summarizes neighborhood preferences as elicited in a US survey study with approximately 1,000 respondents per major ethnic group (Clark and Fossett, 2008). The data show the proportion of respondents per group who found a certain number of co-ethnics ideal in a hypothetical neighborhood containing eight households.

SegregationDiscreteChoice models these preferences with single-peaked neighborhood attractiveness functions. More precisely, for a household with ideal share T , the attractiveness A of a neighborhood with a real share s of co-ethnic households is computed as in Equation (24.1):

$$A(s, T) = \begin{cases} \frac{s}{T} & \text{if } s \leq T \\ 1 - (s - T) & \text{if } s > T \end{cases} \quad (24.1)$$

Figure 24.3 illustrates the substantive assumptions implemented by Equation (24.1). Households find a neighborhood with their ideal composition maximally attractive ($A = 1$). The more a neighborhood deviates from this ideal, the less attractive it is. Further, neighborhoods without

¹⁰ For more information, we refer to our research cluster website where we discuss an extension considering the impact of T on clustering per ethnic group as well as varying the level of density between 70% and 90%: <https://sites.google.com/view/normsandnetworks/addendum-chapter-agent-based-computational-models-in-analytical-sociology>.

¹¹ The model presented in this section was implemented in the separate *NetLogo* program *SegregationDiscreteChoice*.

Table 24.1 Proportions of respondents of four different ethnic groups in the US (African American = Green; Asian = Blue; Hispanic = Yellow; White = Red), choosing an ideal neighborhood with a given number of co-ethnic households (maximum 8). The final column (Modal [T]) indicates the proportion of co-ethnics corresponding with the modal scenario in a category; the left column visualizes a corresponding scenario. The center cell represents the household, and grey and white cells represent same-group and other-group neighbors, respectively. For more details, please see Clark and Fossett (2008, p. 4111)

Tolerance: Number of co-ethnics preferred (out of 8)			Green (20%)	Blue (10%)	Yellow (10%)	Red (60%)	Modal (T)
7/8		Low tolerance (8–6)	11.5	22.6	19.6	22.3	87.5
4/8		Medium tolerance (5–3)	49	51.9	45.9	48.7	50
1/8		High tolerance (2–0)	39.5	25.5	34.5	29	12.5
Total			100	100	100	100	

co-ethnic neighbors are maximally unattractive ($A = 0$), and the higher a household's ideal proportion, the more attractive it deems purely co-ethnic neighborhoods ($A(1, T) = T$).

SegregationDiscreteChoice models households' decision-making following the standard form of discrete choice models as given by Equation (24.2):

$$P_{new} = \frac{e^{\beta A(s_{new}, T)}}{e^{\beta A(s_{new}, T)} + e^{\beta A(s_{old}, T)}} \quad (24.2)^{12}$$

In Equation (24.2), the probability that the new location will be chosen is P_{new} , and the symbols s_{new} and s_{old} indicate the proportion co-ethnics in the potential new and the old neighborhood, respectively. The parameter β scales how closely the choice probability is linked with the relative attractiveness of the two neighborhoods (see procedure *betaDiscreteChoice*). The higher β , the less likely a household chooses the less attractive neighborhood.

The probabilistic discrete choice model can be considered empirically more plausible than the original threshold function. In particular, whether a household leaves a neighborhood for a new one is no longer fully determined by these neighborhoods' compositions, but is also due to random chance. Implicitly, this reflects that in reality households' decisions are also shaped by factors other than neighborhood composition. This model modification illustrates a further important good practice in ABCM.

Good practice #9. Test whether model results of substantive importance are robust if small amounts of random deviations are added to deterministically formulated model assumptions.

¹² In *SegregationDiscreteChoice*, this function is implemented in a mathematically equivalent but computationally more efficient form as $P_{new} = \frac{1}{1 + e^{\beta(A(s_{old}, T) - A(s_{new}, T))}}$.

Figure 24.3 Single-peaked ethnic preference functions for households in three different tolerance categories (based on Table 24.1). Attractiveness reflects how attractive the share of co-ethnics (s) is. Note that a low (high) tolerance means that many (few) co-ethnics are wanted

In *SegregationExtended*, an unhappy household always abandoned a neighborhood and a happy household never did. This deterministic assumption is relaxed in *SegregationDiscreteChoice*. In a real ABCM project, nothing else should be changed when randomness is added to a model for evaluating the robustness of central results. For lack of space, we cannot demonstrate a corresponding test here. However, *SegregationExtended* contains a slider *noise* that allows one to explore whether results change if a small probability of random perturbation of households' happiness is added. Readers are invited to conduct this test and search for explanations if they find that model outcomes change. Generally, testing effects of adding small amounts of random noise is very important in ABCM research as this may not only profoundly change model results, but sometimes also improve the empirical accuracy of model predictions (Macy and Tsvetkova, 2015; Mäs and Helbing, 2020).

The discrete choice model also requires a new method for deciding which households get a chance to relocate. *SegregationDiscreteChoice* adopts the simple solution to select every household with a certain probability per round, given by the parameter *probMoveOption*. Further, the procedure *update-households* selects at random who can move in the next round, and the procedures *move-unhappy-households* and *find-new-spot* now contain code for memorizing the origin location and its attractiveness before the random walk. This allows moving a household back to its original location if it decides to not relocate. Finally, new procedures *attractiveness* and *BinaryChoice* compute $A(s, T)$ for a given neighborhood and conduct the random choice based on Equation (24.2).

The procedure *setup* in *SegregationDiscreteChoice* randomly distributes ethnic preferences T across households based on the relative frequencies of household categories in Table 24.1. Other than in *SegregationExtended*, ethnic preference T (%-similar-wanted) is no longer a global variable but a fixed internal state of an individual household agent (*my-%-similar-wanted*).

3.4.2 Exploring implications

Arguably, *SegregationDiscreteChoice* is more realistic than *SegregationExtended*. But why add all this realism? Manzo (2014) points to an important reason. “Injecting realism” tests whether a model’s ability to provide a generative explanation holds up under further empirically motivated constraints on the mechanism.

For our example, this means we need to test whether *SegregationDiscreteChoice* still meets the research goal and generates segregation higher than would be needed for satisfying households’ ethnic preferences. With single-peaked preferences, this requires new outcome measures. *SegregationDiscreteChoice* computes a set of outcome measures called *oversegregation*, reporting for every household category the average of the difference between the desired share of co-ethnic neighbors (T) and the actual share obtained.

Adding new features to a model also offers researchers the opportunity to formulate and test expectations about effects. In our case, we could ask how a category’s ethnic preference is related to how much it clusters with co-ethnics in the emergent residential distribution. We invite readers to formulate a hypothetical answer based on current model understanding. To test answers, a new measure for spatial *clustering* divides for every household category its average *Percent Similar* by the ethnic group’s population share. Values above 1 indicate that this type of household forms spatial clusters which are ethnically more homogenous than would be expected under random mixing.

Figure 24.4 shows results of a single run (A), and Experiment 3 (B) provides tentative answers. Experiment 3 combines the *world size*, *density*, *radius*, and *ethnic composition* settings of Experiments 1 and 2 with the heterogeneous ethnic preferences based on Table 24.1. Further, *probMoveOption* = 5% approximates moving rates in real cities (Benenson et al., 2009), and $\beta = 12$ sets relative neighborhood attractiveness to be the dominant factor in household’s choices, but leaves room for some random variation.

Figure 24.4A shows a single run charting the evolution of clustering for the three different subcategories of the majority group Red and minority Blue over 2000 rounds. First, we can see that *SegregationDiscreteChoice* still generates the emergence of segregated neighborhoods from initial integration. Except for tolerant majority households, all groups become more clustered than under random mixing. At the same time, the increase of clustering is only steep for medium and high threshold minority groups, while for all other groups clustering increases

Figure 24.4 Panel A: Change of clustering for two selected households (Blue and Red) categories in a single run over 2000 rounds. Panel B: Variance and average oversegregation (10 repetitions) by tolerance category after 1000 rounds

only slightly. Further, we see that tolerance affects residential outcomes. In both ethnic groups, the most tolerant households end up in less segregated neighborhoods than their less tolerant counterparts. However, the precise link between (in)tolerance and clustering is more complex.

Figure 24.4B shows results for the same condition, explicating how ethnic preferences relate to households' average *oversegregation* after 1 000 rounds (Experiment 3). There are only two types of households whose desired neighborhood composition clearly deviates from what they get: tolerant majority households and intolerant minority households. Tolerant majority households have many more co-ethnic neighbors than they would like; intolerant minority households have far too few.

Further refinement of our tentative theory of relocation cascades is required. Based on approximately realistic assumptions about ethnic preferences, the mechanism implemented in *SegregationDiscreteChoice* results only for some types of households in higher segregation than they desire, and for some it even generates less. Again, in a real ABCM research project, further carefully designed experiments would be needed to find out why.

We conclude this example by pointing to an intriguing experiment the reader can conduct with *SegregationDiscreteChoice* in search for an answer. Inspection of single runs with the standard scenario suggests a reason why intolerant minority households fail to become as clustered as they would like. As soon as minority clusters form, vacant spots in them are “invaded” by tolerant majority households seeking to reside in multi-ethnic neighborhoods. Tolerant Red households comprise about 18% of the whole population, more than each of the two minorities, Blue and Yellow. As a consequence, for the taste of intolerant Blue and Yellow households, too many Red households occupy vacant spots in “their” emergent ethnic clusters, driving them away. This explains the relatively low degree of ethnic clustering of intolerant minorities, compared to their medium-tolerant co-ethnics. The explanation can be tested with *SegregationDiscreteChoice*. The model interface comprises a switch *TolerantMajority*, which is set “On” by default. With *TolerantMajority* switched “Off”, all formerly tolerant households in the majority group are turned into medium-tolerant when the model is initialized. Nothing else is changed. A *ceteris-paribus* replication of Experiment 3 with *TolerantMajority* = *Off* confirms the explanation. While the over-segregation index was about -0.5 for these groups in Experiment 3, it rises to about -0.3 and -0.2 for intolerant Blue and Yellow households, respectively. Visual inspection reveals a pattern not visible before. Now, much denser minority clusters emerge, which tend to develop a core-periphery structure in which intolerant households form the center surrounded by more tolerant co-ethnics.

This concluding example underlines two important insights. First, it illustrates (again) that making a model less realistic can be a good practice for better understanding consequences of adding empirically more plausible assumptions. While the empirical data suggest that the household population should contain Red households with highly tolerant preferences, dynamics of the model can be better understood if we remove this category for a moment.

Good practice #10. Move back to analyze the behavior of a model for less realistic conditions if needed to test or develop explanations of model behavior for a more realistic condition.

More central to the program of AS, the example also illustrates how misleading it can be to infer explanations for complex social phenomena from observed individual-level behavior alone. If this were the real world, a researcher could think that strongly segregated core-

periphery neighborhoods of ethnic minorities are caused by these minorities' xenophobic ethnic preferences. As we found, this explanation would be incomplete at best. Minorities with the same preference structure can end up in much less segregated neighborhoods if more majority households in the population are tolerant. Complex interdependencies like this one occur in many social mechanisms. ABCM is a powerful tool to unravel them.

3.5 Even More Realism in Models of Segregation

The simulation experiments shown here illustrate 10 good practices for ABCM. Moreover, they illustrate how ABCM is useful for AS by adding in a step-like fashion realism to a model. Yet, *SegregationDiscreteChoice* could still be deemed a toy-model neglecting numerous factors known to affect empirical residential segregation. It would be misguided to see this as a limitation of ABCM as such. ABCMs of residential segregation were developed that start from models similar to those shown here, incorporating housing prices, income differences, neighborhood reputation (Zhang, 2004; Fossett, 2006), or the real spatial arrangement of a city based on GIS data (Benenson et al., 2009) as further elements.

4. CONCLUSION AND DISCUSSION

ABCM plays a central role in analytical sociology. Our chapter focuses on what we see as its most important strength, the ability to dissect with systematic experimentation and model analysis a potentially complex theoretically proposed social mechanism. This is especially important because full comprehension of an empirically realistic model can be a formidable challenge. Heterogeneity of an agent population, inherent randomness, complex network structures, or local and non-linear interdependencies between multiple agents are but a few of the features many real social systems exhibit. The complexity of social reality is why social scientists increasingly resort to formal modeling and ABCM in the first place (Mäs, Chapter 4, this volume). The gain is that an ABCM is simpler and much more amenable for scientific analysis than social reality itself. Yet, a reasonably empirically plausible model can still be a complex system in itself where the relation between model assumptions and outcomes is far from trivial. A full understanding of why, how, and under which conditions a model is capable of generating outcomes resembling a real social phenomenon is thus a precondition that needs to be met before the ABCM can be considered to provide a valid explanation of the empirical phenomenon it targets.

The good practices we proposed help assure that the strengths of ABCM can be put to its best use for Analytical Sociologists. Our good practices are part of a broader set of established methodologies and principles in ABCM, which are explained and discussed in comprehensive introductions such as (Squazzoni, 2012; Edmonds and Meyer, 2017; Railsback and Grimm, 2019). In addition to model understanding, they also stress the importance of correct implementation of model assumptions (verification), proper use of empirical data for choosing initial conditions and parameter values (calibration), and appropriate testing of model implications against empirical evidence (validation). As the field is moving quickly, interested readers are also advised to consult journals on the intersection of ABCM and sociology, such as *The Journal of Artificial Societies and Social Simulation (JASSS)*, *Computational and Mathematical Organization Theory (CMOT)* or the *Journal of Mathematical Sociology*

(JMS). Readers should further be aware that it has become a generally accepted practice that computer code of published models is made accessible for replication (e.g. CoMSES Net/OpenABM, 2020), following standard protocols such as ODD (Grimm et al., 2020). Debates around the lack of transparency of highly influential models such as those used for predicting mortality and infection rates in the recent COVID-19 crisis have highlighted the special importance of this particular practice (Squazzoni et al., 2020). Finally, one should realize that beyond *NetLogo* there is a multitude of software packages and programming languages available for ABCM. Which of these is the best for a given simulation project depends on many factors, including programming experience of the researchers, available software and hardware, or need for large-scale time-intensive computational experiments. Reviews of the field help to navigate this complex landscape (Abar et al., 2017).

We conclude by reflecting on three debates which are of particular relevance for AS concerning the relation between toy-models and empirically more realistic models, explanation vs. theoretical exploration as research goals, and causal inferences drawn from ABCM.

First, AS rejects the view that empirically unrealistic models should be preferred for considerations of higher elegance or tractability. Yet, this does not mean that model builders have nothing to learn from knowing and understanding “toy-models” available in the literature (Reutlinger et al., 2017). Often, more general mathematical analyses are available or can be constructed for these models, which help to assess and understand whether and under which conditions more detail and realism really change fundamental properties of a model (see, Keijzer et al., 2018, for a recent example). Furthermore, unrealistically simple models can shed light on why seemingly plausible explanations do not work as expected, or highlight new potential explanations in a maximally lucid way, as Schelling’s original model did. Finally, in many cases realistic mechanistic explanations combine elements of several simpler, well-understood mechanisms. For example, while relocation cascades can be one part of the explanation on why segregation arises in a real city, another interconnected mechanism might be increasing opinion polarization between segregated groups, driven by ethnic homophily and social influence processes (Axelrod, 1997; Mäs et al., 2013; Feliciani et al., 2017; Flache, 2018, 2019). Add to this the very real possibility that further processes interfere, such as housing market dynamics, or emergent neighborhood reputations driven by self-reinforcing social influence (van de Rijt, 2019), and it should be clear that researchers are well advised to first get a solid grasp of the dynamics of each of these mechanisms in isolation before attempting to increase realism by integrating them in one model.

Second, a related question is whether ABCM should exclusively be used for providing explanations of empirically observed social phenomena. We believe that this would leave an important potential of ABCM unused. New and fruitful theoretical ideas can emerge from simulation projects asking the “what if” question, focusing on possible but potentially as-yet unobserved consequences of fundamental social processes which are empirically plausible, but are connected in a theoretically new way. ABCM can help to discover and unravel such consequences, which then can give rise to new hypotheses that are testable in experiments or field studies. One example is the insight that a cohesive social network may undermine rather than foster cooperation in the provision of collective goods. This was first highlighted by a purely theoretical ABCM work (Flache and Macy, 1996), which then informed subsequent experimental tests (Flache, 1996; Flache et al., 2017) and resonated later findings from independently conducted field research (Langfred, 1998). Similarly, in studies of opinion polarization, new insights about the role of demographic faultlines were first explored theoretically in computational experiments and then tested in lab experiments (Mäs and Flache, 2013), as was the

insight that complex contagion can play a key role in the diffusion of behaviors in clustered social networks (Centola, 2010), which has inspired further ABCM research showing how incorporation of empirical data on actual social influence processes and network structures can yield new conditions and explain empirical variation in adoption curves between different communities (Manzo et al., 2018). Related to this point, there always is the possibility of computational serendipity. Models designed to explain a specific phenomenon can turn out to generate unexpected implications, which—after a closer look—appear empirically plausible and could inspire new research.

Finally, a key issue in AS is whether an ABCM can actually show that the mechanism it implements really causes the empirical phenomenon it targets. In short, the answer given by analytical sociologists (Hedström and Ylikoski, 2010; Manzo, 2021) is that the more the elements that are constituting a model, e.g. its agents and behavioral and structural assumptions mirroring empirical evidence about them, the more plausible is the claim that the mechanism embodied in the model also causes the phenomenon if the model succeeds in generating macro-outcomes resembling this phenomenon.

Handling the tool of ABCM properly requires adherence to a set of good practices that assure that researchers fully comprehend the complexity of their own theoretical creations. Under this precondition, ABCM is a unique method for providing precise, causal and empirically plausible mechanistic explanations of puzzling phenomena in the social world.

REFERENCES

- Abar, S., G.K. Theodoropoulos, P. Lemarinier and G.M.P. O'Hare (2017), "Agent based modelling and simulation tools: A review of the state-of-art software", *Computer Science Review*, **24**: 13–33. doi: 10.1016/j.cosrev.2017.03.001.
- Axelrod, R. (1997), "The dissemination of culture: A model with local convergence and global polarization", *Journal of Conflict Resolution*, **41** (2): 203–26. doi: 10.1177/0022002797041002001.
- Baldassarri, D. and P. Bearman (2007), "Dynamics of political polarization", *American Sociological Review*, **72** (5): 784–811. doi: 10.1177/000312240707200507.
- Bendor, J. and P. Swistak (2001), "The evolution of norms", *American Journal of Sociology*, **10** (6): 1493–545. doi: 10.1086/321298.
- Benenson, I., E. Hatna and E. Or (2009), "From Schelling to spatially explicit modeling of urban ethnic and economic residential dynamics", *Sociological Methods and Research*, **37** (4): 463–97. doi: 10.1177/0049124109334792.
- Bianchi, F., A. Flache and F. Squazzoni (2020), "Solidarity in collaboration networks when everyone competes for the strongest partner: A stochastic actor-based simulation model", *Journal of Mathematical Sociology*, **44** (4): 249–66. doi: 10.1080/0022250X.2019.1704284.
- Bianchi, F. and F. Squazzoni (2015), "Agent-based models in sociology", *Wiley Interdisciplinary Reviews: Computational Statistics*, **7** (4): 284–306. doi: 10.1002/wics.1356.
- Breen, R. (2009), "Game theory", in P. Hedström and P. Bearman (eds), *The Oxford Handbook of Analytical Sociology*, Oxford: Oxford University Press, pp. 619–38.
- Bruch, E.E. and J. Atwell (2015), "Agent-based models in empirical social research", *Sociological Methods and Research*, **44** (2): 186–221. doi: 10.1177/0049124113506405.
- Bruch, E.E. and R.D. Mare (2006), "Neighborhood choice and neighborhood change", *American Journal of Sociology*, **112** (3): 667–709. doi: 10.1086/507856.
- Bruch, E.E. and R.D. Mare (2009), "Segregation dynamics", in P. Hedström and P. Bearman (eds), *The Oxford Handbook of Analytical Sociology*, Oxford: Oxford University Press, pp. 269–93. doi: 10.1093/oxfordhb/9780199215362.013.12.
- Castellano, C., S. Fortunato and V. Loreto (2009), "Statistical physics of social dynamics", *Reviews of Modern Physics*, **81** (2): 591–646. doi: 10.1103/RevModPhys.81.591.
- Centola, D. (2010), "The spread of behavior in an online social network experiment", *Science*, **329** (5996): 1194–7. doi: 10.1126/science.1185231.
- Centola, D., J.C. González-Avella, V.M. Eguíluz and M. San Miguel. (2007), "Homophily, cultural drift, and the co-evolution of cultural groups", *Journal of Conflict Resolution*, **51** (6): 905–29. doi: 10.1177/0022002707307632.

- Clark, W.A.V. and M. Fossett (2008), "Understanding the social context of the Schelling segregation model", *Proceedings of the National Academy of Sciences of the United States of America*, **105** (11): 4109–14. doi: 10.1073/pnas.0708155105.
- CoMSES Net/OpenABM (2020), *Computational Model Library*. <https://www.comses.net/>.
- Edmonds, B. and R. Meyer (2017), *Simulating Social Complexity*, *Simulating Social Complexity*, Berlin, Heidelberg: Springer. doi: 10.1007/978-3-540-93813-2.
- Epstein, J.M. (2006), *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton, NJ: Princeton University Press. doi: 10.5038/2162-4593.11.1.8.
- Feliciani, T., A. Flache and J. Tolsma (2017), "How, when and where can spatial segregation induce opinion polarization? Two competing models", *Journal of Artificial Societies and Social Simulation*, **20** (2): 6. doi: 10.18564/jasss.3419.
- Flache, A. (1996), *The Double Edge of Networks. An Analysis of the Effect of Informal Networks on Cooperation in Social Dilemmas*, Amsterdam: Thesis Publishers.
- Flache, A. (2018), "About renegades and outgroup haters: Modeling the link between social influence and intergroup attitudes", *Advances in Complex Systems*, **21** (6–7): 1850017. doi: 10.1142/S0219525918500170.
- Flache, A. (2019), "Social Integration in a Diverse Society: Social Complexity Models of the Link between Segregation and Opinion Polarization", in F. Abergel, B.K. Chakrabarti, A. Chakraborti, N. Deo and K. Sharma (eds), *New Perspectives and Challenges in Econophysics and Sociophysics (New Economic Windows)*, Springer, Cham, pp. 213–28. doi: 10.1007/978-3-030-11364-3_15.
- Flache, A., D.M. Bakker, M. Mäs and J. Dijkstra et al. (2017), "The Double Edge of Counter-Sanctions. Is Peer Sanctioning Robust to Counter-Punishment but Vulnerable to Counter-Reward?", in B. Jann and W. Przepiorka (eds), *Social Dilemmas, Institutions, and the Evolution of Cooperation*, Berlin: De Gruyter Oldenbourg, pp. 280–301. doi: 10.1515/9783110472974-014.
- Flache, A., M. Mäs, T. Feliciani, E. Chattoe-Brown, G. Deffuant, S. Huet and J. Lorenz (2017), "Models of social influence: Towards the next frontiers", *Journal of Artificial Societies and Social Simulation*, **20** (4): 2. doi: 10.18564/jasss.3521.
- Flache, A. and M.W. Macy (1996), "The weakness of strong ties: Collective action failure in a highly cohesive group", *Journal of Mathematical Sociology*, **21** (1–2): 3–28. doi: 10.1080/0022250X.1996.9990172.
- Flache, A. and M.W. Macy (2011), "Small worlds and cultural polarization", *Journal of Mathematical Sociology*, **35** (1–3): 146–76. doi: 10.1080/0022250X.2010.532261.
- Fossett, M. (2006), "Including preference and social distance dynamics in multi-factor theories of segregation", *Journal of Mathematical Sociology*, **30** (3–4): 289–98. doi: 10.1080/00222500500544151.
- Goldstein, D.G. (2009), "Heuristics", in P. Hedström and P. Bearman (eds), *The Oxford Handbook of Analytical Sociology*, Oxford: Oxford University Press, pp. 140–67. doi: 10.1093/oxfordhb/9780199215362.013.7.
- Grimm, V., S.F. Railsback, C.E. Vincenot, U. Berger, C. Gallagher, D.L. DeAngelis, B. Edmonds, J. Ge, J. Giske, J. Groeneveld, A.S.A. Johnston, A. Milles, J. Nabe-Nielsen, J.G. Polhill, V. Radchuk, M.-S. Rohwäder, R.A. Stillman, J.C. Thiele and D.I. Ayllón (2020), "The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism", *Journal of Artificial Societies and Social Simulation*, **23** (2): 7. doi: 10.18564/jasss.4259.
- Hedström, P. and P. Bearman (2009), "What Is Analytical Sociology All About? An Introductory Essay", in P. Hedström and P. Bearman (eds), *The Oxford Handbook of Analytical Sociology*, Oxford: Oxford University Press, pp. 3–24.
- Hedström, P. and P. Ylikoski (2010), "Causal mechanisms in the social sciences", *Annual Review of Sociology*, **36** (1): 49–67. doi: 10.1146/annurev.soc.012809.102632.
- Hegselmann, R. (2017), "Thomas C. Schelling and James M. Sakoda: The intellectual, technical, and social history of a model", *Journal of Artificial Societies and Social Simulation*, **20** (3): 15. doi: 10.18564/jasss.3511.
- Hegselmann, R. and A. Flache (1998), "Understanding complex social dynamics: A plea for cellular automata based modelling", *Journal of Artificial Societies and Social Simulation*, **1** (3). <http://www.soc.surrey.ac.uk/JASSS/1/3/1.html>.
- Holm, S., R. Lemm, O. Thees and L.M. Hilty (2016), "Enhancing agent-based models with discrete choice experiments", *Journal of Artificial Societies and Social Simulation*, **19** (3): 3. doi: 10.18564/jasss.3121.
- Keijzer, M.A., M. Mäs and A. Flache (2018), "Communication in online social networks fosters cultural isolation", *Complexity*, 9502872. doi: 10.1155/2018/9502872.
- Keuschnigg, M., N. Lovsjö and P. Hedström (2018), "Analytical sociology and computational social science", *Journal of Computational Social Science*, **1** (1): 3–14. doi: 10.1007/s42001-017-0006-5.
- Langfred, C.W. (1998), "Is group cohesiveness a double-edged sword?", *Small Group Research*, **29** (1): 124–43. <https://doi.org/10.1177/1046496498291005>.
- Macy, M.W. and A. Flache (2002), "Learning dynamics in social dilemmas", *Proceedings of the National Academy of Sciences of the United States of America*, **99** (suppl. 3): 7229–36. doi: 10.1073/pnas.092080099.
- Macy, M.W. and A. Flache (2009), "Social Dynamics from the Bottom Up: Agent-based Models of Social Interaction", in P. Hedström and P. Bearman (eds), *The Oxford Handbook of Analytical Sociology*, Oxford: Oxford University Press, pp. 245–68. doi: 10.1093/oxfordhb/9780199215362.013.11.

- Macy, M.W. and M. Tsvetkova (2015), "The signal importance of noise", *Sociological Methods and Research*, **44** (2): 306–28. doi: 10.1177/0049124113508093.
- Manzo, G. (2014), "Data, Generative Models, and Mechanisms: More on the Principles of Analytical Sociology", in G. Manzo (ed.), *Analytical Sociology: Actions and Networks*, Chichester, UK: Wiley pp.4–52. doi: 10.1002/9781118762707.ch01.
- Manzo, G. (2021), *Agent-based Models and Causal Inference*, Hoboken, NJ: Wiley.
- Manzo, G., S. Gabbrellini, V. Roux and F.N. M'mbogori (2018), "Complex contagions and the diffusion of innovations: Evidence from a small-N study", *Journal of Archaeological Method and Theory*, **25** (4): 1109–54. doi: 10.1007/s10816-018-9393-z.
- Mark, N.P., L. Smith-Lovin and C.L. Ridgeway (2009), "Why do nominal characteristics acquire status value? A minimal explanation for status construction", *American Journal of Sociology*, **115** (3): 832–62. doi: 10.1086/606142.
- Mäs, M. (2021), "Analytical Sociology and Complexity Research", in G. Manzo (ed.), *Research Handbook on Analytical Sociology*, Cheltenham, UK: Edward Elgar.
- Mäs, M. and A. Flache (2013), "Differentiation without distancing. Explaining bi-polarization of opinions without negative influence", *PLoS ONE*, **8** (11), p. e74516. doi: 10.1371/journal.pone.0074516.
- Mäs, M., A. Flache, K. Takács and K.A. Jehn (2013), "In the short term we divide, in the long term we unite: Demographic crisscrossing and the effects of faultlines on subgroup polarization", *Organization Science*, **24** (3): 716–36. doi: 10.1287/orsc.1120.0767.
- Mäs, M. and D. Helbing (2020), "Random deviations improve micro–macro predictions: An empirical test", *Sociological Methods and Research*, **49** (2): 387–417. doi: 10.1177/0049124117729708.
- Mäs, M. and K.-D. Opp (2016), "When is ignorance bliss? Disclosing true information and cascades of norm violation in networks", *Social Networks*, **47**: 116–29. doi: 10.1016/j.socnet.2016.05.004.
- Railsback, S.F. and V. Grimm (2019), *Agent-Based and Individual-Based Modeling: A Practical Introduction*, 2nd ed, Princeton, NJ: Princeton University Press.
- Reutlinger, A., D. Hangleiter and S. Hartmann (2017), "Understanding (with) toy models", *The British Journal for the Philosophy of Science*, **69** (4): 1069–99. doi: 10.1093/bjps/axx005.
- van de Rijt, A. (2019), "Self-correcting dynamics in social influence processes", *American Journal of Sociology*, **124** (5): 1468–95. doi: 10.1086/702899.
- van de Rijt, A., D. Siegel and M.W. Macy (2009), "Commentary and Debate", *American Journal of Sociology*, **114** (4): 1166–80. doi: 10.1086/588795.
- van Gent, W., M. Das and S. Musterd (2019), "Sociocultural, economic and ethnic homogeneity in residential mobility and spatial sorting among couples", *Environment and Planning A*, **51** (4): 891–912. doi: 10.1177/0308518X18823754.
- Sakoda, J.M. (1971), "The checkerboard model of social interaction", *The Journal of Mathematical Sociology*, **1** (1): 119–32. doi: 10.1080/0022250X.1971.9989791.
- Schelling, T.C. (1971), "Dynamic models of segregation", *The Journal of Mathematical Sociology*, **1** (2): 143–86. doi: 10.1080/0022250X.1971.9989794.
- Simon, H.A. (1982), *Models of Bounded Rationality*, Cambridge, MA: MIT Press. doi: 10.7551/mitpress/4711.001.0001.
- Squazzoni, F. (2012), *Agent-Based Computational Sociology*, Chichester, UK: Wiley. <https://doi.org/10.1002/9781119954200>.
- Squazzoni, F., J.G. Polhill, B. Edmonds, P. Ahrweiler, P. Antosz, G. Scholz, É. Chappin, M. Borit, H. Verhagen, F. Giardini and N. Gilbert (2020), "Computational models that matter during a global pandemic outbreak: A call to action", *Journal of Artificial Societies and Social Simulation*, **23** (2): 10. doi: 10.18564/jasss.4298.
- Wilensky, U. (1997), "NetLogo Segregation model", Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. <http://ccl.northwestern.edu/netlogo/models/Segregation>.
- Wilensky, U. (1999), "NetLogo". Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University. <http://ccl.northwestern.edu/netlogo/>.
- Wilensky, U. and W. Rand (2015), *An Introduction to Agent-Based Modeling: Modeling Natural, Social and Engineered Complex Systems with NetLogo*, Cambridge, MA: MIT Press.
- Witek, R.P.M., T.A.B. Snijders and V. Nee (2013), "Introduction: Rational Choice Social Research", in *The Handbook of Rational Choice Social Research*, Stanford, CA: Stanford University Press, pp. 1–30.
- Wooldridge, M. and N.R. Jennings (1995), "Intelligent agents: Theory and practice", *The Knowledge Engineering Review*, **10** (2): 115–52. doi: 10.1017/S0269888900008122.
- Xie, Y. and X. Zhou (2012), "Modeling individual-level heterogeneity in racial residential segregation", *Proceedings of the National Academy of Sciences*, **109** (29): 11646–51. doi: 10.1073/PNAS.1202218109.
- Young, H.P. (1998), *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*, Princeton, NJ: Princeton University Press.
- Zhang, J. (2004), "A dynamic model of residential segregation", *Journal of Mathematical Sociology*, **28** (3): 147–70. doi: 10.1080/00222500490480202.
- Zschache, J. (2018), "Melioration learning in iterated public goods games: The impact of exploratory noise", *The Journal of Mathematical Sociology*, **42** (1): 1–16. doi: 10.1080/0022250X.2017.1396983.