

Agent-Based Models as Etio-Prognostic Explanations

Olaf Dammann

Tufts University

Abstract

Agent-based models (ABMs) are one type of simulation model used in the context of the COVID-19 pandemic. In contrast to equation-based models, ABMs are algorithms that use individual agents and attribute changing characteristics to each one, multiple times during multiple iterations over time. This paper focuses on three philosophical aspects of ABMs as models of causal mechanisms, as generators of emergent phenomena, and as providers of explanation. Based on my discussion, I conclude that while ABMs cannot help much with causal inference, they can be viewed as etio-prognostic explanations of illness occurrence and outcome.

Keywords: Explanation, Causation, Simulation, Modelling, COVID-19.

1. Introduction

Computational modeling and simulation of real-life scenarios have become a mainstay in health research and the biosciences. My goal in this interdisciplinary paper, written from my personal perspective as physician, epidemiologist, and philosopher, is to provide an analysis of the explanatory scope of agent-based models (ABMs), one particular kind of modeling technique employed in the context of the COVID-19 pandemic (Silva et al. 2020, Cuevas 2020, Truskowska et al. 2021, Shamil et al. 2021, Hoertel et al. 2020, Staffini et al. 2021, Kerr et al. 2021). It is not my intention to review these papers in detail; suffice it to say that they are all part of the general endeavor to tackle important population health problems posed by the COVID pandemic and have made considerable contributions to our understanding of epidemiological dynamics of this global health crisis. Instead, my discussion will focus on three philosophical aspects of ABMs as models of causal mechanisms, generators of emergent phenomena, and providers of explanation.

I will start by introducing ABMs and why they are generally considered helpful (§2). Part of their epistemological value is that they are thought to provide explanations of biological and social mechanisms (§3). One account of ABMs, featured prominently on the Columbia School of Public Health website, has ABMs as models of causal mechanisms of interactions of characteristics that may include

impossible or unethical connections (§4) and that generate emergent phenomena (§5). The paper ends with the proposal to consider ABMs as helpful in generating etio-prognostic explanations (§6).

2. Agent-Based Models

As any other computational model, an ABM is an algorithm with inputs, computations, and outputs. In public health, ABMs are generally conceptualized as “a computational approach in which agents with a specified set of characteristics interact with each other and with their environment according to predefined rules” (Tracy, Cerdá, and Keyes 2018: 77). What exactly does that mean and why should this be helpful?

2.1 What Are ABMs?

An ABM (sometimes also called individual-based model or IBM) is a computer program that simulates changes in populations over time based on the ‘behavior’ of ‘agents’ who have a set of characteristics and ‘interact’ in predefined and stochastically modeled ways. This kind of simulation is often called *microsimulation* because phenomena are modeled at the micro-level (the individual agent) and results are observed at the macro-level, the level of the simulated population. Starting values and conditions for transition of agents from one state to another (for example, from non-infected to infected or from alive to dead) are defined by the programmer. Running the program will result in iterations of changes in these conditions over time. Ending conditions at the macro-level are the outcome of the model. Since the attribution of particular values to individual agents is done by randomly allocating values selected from a probability distribution with set constraints, each run of the algorithm will result in a different outcome. Multiple, oftentimes many runs need to be performed to arrive at a range of outcomes that defines an outcome distribution. The results of ABMs are non-deterministic such as those of equation-based models (EBMs). For a comparison of ABMs and EBMs, see (Van Dyke Parunak, Savit, and Riolo 1998).

Agent-based modeling is frequently used in theoretical infectious disease epidemiology (Venkatramanan et al. 2018). As outlined by Hunter and colleagues, ABMs are considered superior to EBMs (such as those that generate the now very familiar COVID-19 incidence and mortality curves) because they allow for the modelling of the behavior of individuals based on social interaction rules and a probabilistic attribution of such behaviors to the agents in a model (Hunter, Mac Namee, and Kelleher 2017). Agent-based models have to consider four major related aspects: disease, society, movement, and environment. They have to model disease-specific conditions of occurrence and duration, characteristics of the society (population) and how its members move through virtual space and interact with one another in the environment that population is situated in. The result is a highly complex representation of how population parameters change over time with regard to, e.g., disease incidence or mortality rates. Let me note right here that ABMs involve equations as well. However, the underlying equations let a set of variables undergo iterative changes over a pre-defined timeframe so that such changes over time *in each agent* contribute to an overall change at the population level.

Let us go through the published description of one ABM-based microsimulation and parse out its individual observational and inferential components.

2.2 Example: ABM of the COVID-19 Epidemic in France

Hoertel and coworkers published

a stochastic agent-based microsimulation model of the COVID-19 epidemic in France. [They] examined the potential impact of post-lockdown measures, including physical distancing, mask-wearing and shielding individuals who are the most vulnerable to severe COVID-19 infection, on cumulative disease incidence and mortality, and on intensive care unit (ICU)-bed occupancy. While lockdown is effective in containing the viral spread, once lifted, regardless of duration, it would be unlikely to prevent a rebound. Both physical distancing and mask-wearing, although effective in slowing the epidemic and in reducing mortality, would also be ineffective in ultimately preventing ICUs from becoming overwhelmed and a subsequent second lockdown. However, these measures coupled with the shielding of vulnerable people would be associated with better outcomes, including lower mortality and maintaining an adequate ICU capacity to prevent a second lockdown (Hoertel et al. 2020: 1417).

The goal of the model was to simulate the *effect* of changing measures after the first lockdown in France such as social distancing, mask-wearing, and shielding the most vulnerable. Outcomes measures (variables) at the *population* level (macro-level) were a rebound, second lockdown, epidemic slow down, intensive care unit admission rates, mortality, as well as combinations of the above. In order to arrive at their results, investigators needed to model events at the *individual* level (micro-level), including

194 parameters related to French population characteristics (n = 140), social contacts (n = 33) and SARS-CoV-2 characteristics (n = 21) [...]. Parameter values on population characteristics were based on data from the French National Statistical Institute (INSEE) and Santé Publique France. Parameters related to social contacts were based on prior studies (n = 11) or assumptions when no data were available (n = 22). Finally, parameters on disease characteristics were based on data from Institut Pasteur and London Imperial College, except for two unknown key parameters of the epidemic: contamination risk and proportion of undiagnosed COVID-19 cases, which were simultaneously estimated through model calibration (Hoertel et al. 2020) (quote from online material available at <https://www.nature.com/articles/s41591-020-1001-6#Sec2>; accessed 06/13/2021).

In essence, almost two hundred individual and population characteristics were modeled over time and the resulting changes at the population level were observed. Circling back to my tripartite goal in this paper to explore ABMs as (a) models of mechanisms, (b) generators of emergent phenomena, and (c) providers of explanation, the (a) mechanisms would be the joint changes over time among the agents of the ABM that (b) lead to certain population-based emergent phenomena, and (c) observing the model values change and results emerge would provide an explanation. The central question I ask in this paper is, an explanation of *what* exactly this might be.

2.3 Why Are ABMs Considered Helpful?

Obviously, ABMs are created for a purpose. In the general context of our current discussion, any modeler of an epidemic (pandemics included) has at least three goals. First, they want to *understand the dynamics* of the epidemic in terms of background conditions of population and environment. Thus, the first goal is to find a *causal-mechanical explanation* of why and how the infection spreads in populations. Second, modelers want to create an algorithm that allows them to *predict* how the epidemic will evolve over time. Third, modelers want to explore changes in model outcomes in response to parameter changes. In an iterative fashion, the algorithm can be modified to get closer and closer to predictions that can be confirmed or rejected by real-life data as time goes by.

I have mentioned above that one of the motivations to create ABMs is that they are considered superior to EBMs in terms of being more realistic (Hunter, Mac Namee, and Kelleher 2017). Equation-based models are simple, static, and deterministic, because they are built like a mathematical formula such as a regression equation that gives a result on a dependent variable based on the value of one or more independent variables. Once the regression equation is derived from an observational study in a certain population, any new observation can be plugged into the regression formula and a predicted value for the dependent variable can be obtained. They are static in the sense of being non-dynamic. This means that once a regression equation is created, it doesn't change. If one wants to look at other combinations of variables, different starting conditions, or changes over time, one needs to create new equations. And they are deterministic, because the value of the dependent variable is fixed once the values of the independent variables are fixed. There is not much room for "natural variation" in equation-based modelling.

Consider the following excerpt from an outline of ABMs on the website of one of the major schools of public health in the United States:

Agent-based models are computer simulations used to study the interactions between people, things, places, and time. They are stochastic models built from the bottom up meaning individual agents (often people in epidemiology) are assigned certain attributes. The agents are programmed to behave and interact with other agents and the environment in certain ways. These interactions produce emergent effects that may differ from effects of individual agents. Agent-based modeling differs from traditional, regression-based methods in that, like systems dynamics modeling, it allows for the exploration of complex systems that display non-independence of individuals and feedback loops in causal mechanisms. It is not limited to observed data and can be used to model the counterfactual or experiments that may be impossible or unethical to conduct in the real world (<https://www.publichealth.columbia.edu/research/population-health-methods/agent-based-modeling>), accessed 06-04-2021).

Let me henceforth refer to this blurb as the *Columbia account of ABM* and rephrase its elements as three epistemological statements we can use as a guideline for the next sections of this paper.

Agent-based models are epistemologically helpful because they

1. enable the exploration of complex systems characterized by (among other things) non-independence of individuals and feedback loops in *causal mechanisms*, i.e., the sequential processes of changes in agent "behavior" that

connect the initial states among agents and outcomes established at the population level;

2. support the study of interactions at the levels of people, things, places, and time between programmed behaviors of and *interactions* between agents that produce *emergent effects*;
3. can explore mechanisms in ways that are *impossible* in observational and experimental research.

I will now turn to each one of these three epistemological benefits of ABMs in §3-5, respectively.

3. Mechanisms and Causes

3.1 Biological Mechanisms

In the basic biosciences, mechanistic views of biological processes appear to include the notions of *action* and *behavior* when it comes to the observation of changes among the mechanism's components and also the changes that occur as part of the result of the process. For example, Olaf Wolkenhauer writes that systems biologists are interested in finding out "how biological function emerges from the interactions between the components of living systems and how these emergent properties enable and constrain the behavior of those components" (Wolkenhauer 2014). First, note that Wolkenhauer says that biological function "emerges" from the interactions of components. We will come back to emergence in the next section. Second, consider this account of systems biology in light of one of the more frequently cited definitions of a mechanism in philosophy of science: "Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (Machamer, Darden, and Craver 2000). Taken together, the two accounts allow for the inference that at least some systems biologists see their work as identifying biological *mechanisms*.

Wolkenhauer confirms this by saying that "[t]he iterative cycle of data-driven modeling and model-driven experimentation [...] helps in identifying new mechanistic details of cell-biological processes and previously unidentified regulatory *interactions* in the system" (italics mine). Thus, another important similarity between Wolkenhauer's account of computational systems biology and Machamer, Darden, and Craver's account of mechanism is that both refer to some sort of *action*, as in "interactions" and "activities", suggesting that at least some bioscientists think that biological mechanisms are characterized by interactions and activities among the element of those mechanisms.

Let us now move from biological to population mechanics. It seems that population health scientists have a similarly mechanistic view of population health as biologists view biological processes as mechanisms. Consider, for example, this quote from the book "The Future of the Public's Health in the 21st Century" published by the Institute of Medicine (U.S.) Committee on Assuring the Health of the Public in the 21st Century (Medicine 2003): "(a)spects of discrimination might influence health through any number of mechanisms, including (socio-economic status)" (61) and "[t]here are several plausible mechanisms by which social cohesion might influence health through contextual effects" (71). These quotes raise the question how *social mechanisms* are conceptualized.

3.2 Social Mechanisms

Let us first consider who or what the elements of social mechanisms are. Stinchcombe suggests that “[m]echanisms in a theory are defined here as bits of theory about entities at a different level (e.g., individuals) than the main entities being theorized about (e.g., groups), which serve to make the higher-level theory more supple, more accurate, or more general” (Stinchcombe 1991). For our present discussion of epidemic ABMs as models of social mechanisms, the agent would be a representation of an individual person and the entirety of agents would be a representation of a social group or population. From this perspective, individuals (represented by agents in ABMs) would be the actors. In what might be the most frequently cited text on social mechanisms *as explanations*, Hedström and Svedberg (1998) confirm this when they state that their concept of social mechanism is based on four core principles, i.e., action, precision, abstraction and reduction (Hedström and Swedberg 1998). They write that

[t]he first of these principles—explanations based on actions—means, among other things, that it is actors and not variables who do the acting. A mechanism-based explanation is not built upon mere associations between variables but always refers directly to causes and consequences of individual action oriented to the behavior of others (*ibid.*).

Are ABMs, therefore, models of social mechanisms? The following quote seems to answer in the affirmative. Conte and Paolucci write that

[a] generative explanation of an observed social phenomenon consists of describing it in terms of the external (environmental and social) and internal (behavioral) mechanisms that generate them, rather than by inferring causes from observed covariations. This is a vital property of explanation, which cannot easily be realized otherwise. When describing agent behavior by means of other formalisms (logic-based or numeric), we describe behavior from the outside, as perceived by an observer, but do not describe the way it is generated. ABM explains (sic) behavior from within, in terms of the mechanisms that are supposed to have generated it, that is, the mechanisms that operate in the agent when s/he behaves one way or another (Conte and Paolucci 2014).

However, note that Conte and Paolucci carefully distinguish between mechanisms as *natural constituents* of the real processes the ABM is supposed to be a model of, and the *structural and functional blueprint* for agents’ interactions coded into the model algorithm. They appear to see social phenomena as generated (produced) by mechanisms (external and internal) and the advantage of ABMs over other kinds of models as their capability to offer a mechanistic explanation of system behavior. Topping and colleagues make it eminently clear that the mechanisms are *built into the model*. They begin their article (about their ecological ABM model of the European brown hare) as follows:

Agent-based models (ABMs) are gaining popularity in most scientific fields due to their ability to describe complex systems from first principles. Yet, they are also criticised for being ‘black boxes’ and impossible to fully understand. This is mainly due to the difficulty of testing, documenting and communicating the wealth of mechanisms built into such models (Topping, Høye, and Olesen 2010).

This view is confirmed by a group of researchers who designed an ABM on social distancing, testing, contact tracing, and quarantine on the occurrence of SARS-CoV-19 infections. Referring to multiple scenarios they modeled they write that “[t]he above scenarios are mechanistically simulated on the multi-layer network [...] by allowing different interactions (between effective contacts) according to the simulated strategy” (Aleta et al. 2020). Clearly, this team stresses the point that the *simulation* is mechanistic. They do *not* say that they think that the real-life phenomena they model are mechanistic as well. However, what other reason could they have creating mechanistic models than being convinced that the modeled social and behavioral processes are mechanistic as well? Perhaps, we can paraphrase Nancy Cartwright’s “no causes in, no causes out” here as “no mechanisms in, no mechanisms out” (Cartwright 1989), meaning that only if we already have mechanistic background information can we see ABMs as mechanisms. If ABMs are considered mechanistic explanations of a certain phenomenon, they explain the occurrence of the phenomenon as resulting from a mechanism by demonstrating that the phenomenon does indeed occur because of the mechanism modeled by the ABM. However, this does not yet allow for the inference that the phenomenon must be due to this mechanism. To do that, other potential mechanisms and the possibility of chance need to be ruled out, and of course the existence of the mechanism needs to be demonstrated by real world data.

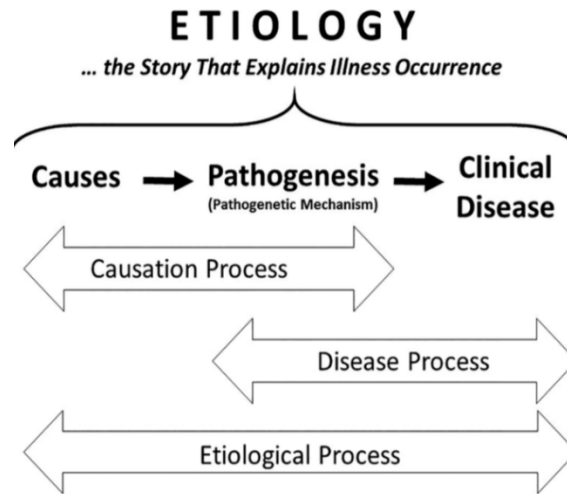
3.3 Causal Mechanisms

Until now, I have tried to avoid the topic of causality because I wanted my focus to be on the role of ABMs as explaining mechanisms without reference to causation. However, some modelers talk about *causal mechanisms* when they talk about the relation between how they see causality in the world and in their models. Consider the Columbia account of ABM above: “[...] exploration of complex systems that display non-independence of individuals and feedback loops in *causal mechanisms*” (italics mine). This notion resonates with Tracy and coworkers’ view that

ABMs are well suited to the exploration of causal mechanisms given their ability to incorporate multiple interacting causes and to test competing theories about causation, thus further elucidating what we do and do not know about how a given outcome arises (Tracy, Cerdá, and Keyes 2018: 85).

It almost seems as if knowledge about mechanisms is considered crucial because it can provide knowledge about causation. As much as I agree that causes and mechanisms have a very close working relationship, they are two very different things. Indeed, Dammann has argued for a distinction between causes and mechanisms in the context of illness occurrence as separate, but closely related compo-

Figure 1. The etiological stance conceptualizes disease occurrence as a process. The first phase (causation process) includes causes and the subsequent pathogenetic mechanism they induce. The second phase (disease process) includes the pathogenesis and clinical disease. Knowledge about both (etiological process), combined with knowledge about the action of other contributors to the etiological process at all of its levels, can provide useful etiological explanations (reprinted with permission from Dammann 2017).



nents of the *causation process* that represents the initial phase of illness etiology (Figure 1) (Dammann 2017). According to this account, causes initiate mechanisms that in turn culminate in clinical illness. Within this etiological scenario *all* mechanisms are causal because they link causes and their outcomes. However, this does not necessarily mean that all mechanisms must be causal. If non-causal mechanisms exist, and if ABMs can model *any* kind of mechanism, then not all mechanisms that can be represented in ABMs are causal. Therefore, any method that is supposed to extract information about *causal* mechanisms from ABMs would need to distinguish between causal and non-causal mechanisms in ABMs. On the other hand, it could be that *all* mechanisms are causal, simpliciter. We would not need to distinguish between causal and non-causal mechanisms because the latter do not exist. If all mechanisms are causal, and ABMs can model any mechanism, ABMs could be used in the exercise of generating causal-mechanical (etiological) explanations. If not, we would, again, need criteria for separating causal from non-causal mechanisms.

What could non-causal mechanisms look like apart from, say, non-functional mechanisms such as repetitive loops in which model parameters do not change? I am referring back to Machamer, Darden, and Craver's definition of mechanism as "entities and activities organized such that they are productive of regular changes". I take this to mean that mechanisms *produce* change. Mechanisms are the way by which causes make a difference. From this perspective, it would seem that all mechanisms are causal. Therefore, if ABMs represent mechanisms, and if all mechanisms are causal, then ABMs are representations of causal mechanisms. Does this mean that ABMs can be used as tools in causal inference?

3.4 ABMs and Causal Inference in Epidemiology

Let us assume that ABMs include causal interactions *by definition*. They are programmed to reflect a causal relationship between variables whenever one is coded to change in response to another. Indeed, this is a representation of a common causal intuition: X causes Y if Y changes whenever X changes. (I sometimes call this, somewhat informally, the light switch intuition.) It includes traditional philosophical notions of causation as regularity, difference-making, dependence, and so forth. However, I see the argument that ABMs are helpful in causal inference as being based on circular reasoning: causality in, causality out (paraphrasing Cartwright, again). ABMs cannot help with causal inference because inference is the bottom-up support of a proposition by observed data. ABMs cannot provide such data because the data they provide are top-down, generated computationally by algorithms. Yes, the *model of the algorithm* itself, e.g., the assumptions and almost 200 parameters used by Hoertel et al in our COVID-19 epidemic example above, may be based on observed data (such as disease incidence, contact frequency among agents, etc.), but the algorithm is designed to produce a result. Thus, the result is caused by the algorithm, and that causal fact does not support the notion that the underlying observed data are reflective of a causal scenario, but only the notion that the algorithm functions as a causal mechanism, and that an algorithmic causal mechanism can be interpreted as a depiction of an envisioned causal mechanism in real-life, but not as evidence supporting the inference that the modeled real-life scenario is causal or the inference that a real-life causal even exists. An algorithmic causal mechanism only shows that such mechanism has the potential to yield the modeled phenomenon. The epistemic gain is demonstrative in a theoretical way (*in silico*), but not in a practical way as in experimentation with animal models (*in vivo*). Both *in silico* and *in vivo* demonstrations confirm the possibility of a role for the mechanism in the purported etiological process, but they do not confirm that it does indeed play that role in real life scenarios.

Another caveat comes from the observation that those who argue for or against methods for causal inference via some method or another usually do so while depending on their own, implicit and often unstated intuitions about the nature of causality (Casini and Manzo 2016). What are epidemiologists' definitions of "causation"? Susser simply states that "a cause is what makes a difference" (Susser 1991). A classic paper on the counterfactual definition of *causal effect* in epidemiology includes my favorite statement "in ideal randomized experiments, association *is* causation" (Hernán 2004). My problem with this paper is, however, that it contrasts the term *causal effect* with the term *effect* because the latter is, according to the author, commonly used to mean "simply statistical association". I think the term causal effect *introduces confusion*, because there is simply no such thing as a non-causal effect. All effects are results of causal mechanisms *by definition*, although the exact mechanism itself is not always known. The more important issue here is, however, that Hernán sees the (population) definition of causal effect simply as tied to a probability differential of developing an outcome under two different exposure conditions (yes or no):

We define the probability $\Pr[Y_a=1]$ as the proportion of subjects that would have developed the outcome Y had all subjects in the population of interest received

exposure value a . We also refer to $\Pr[Y_a=1]$ as the risk of Y_a . The exposure has a causal effect in the population if $\Pr[Y_{a=1}=1] \neq \Pr[Y_{a=0}=1]$ (Hernán 2004: 266).

This definition strikes me as applicable to “statistical association”, but by no means would I subscribe to the view that it defines “causal effect” without further explication of what Hernán *means* by “causal effect”. Unless he intends to suggest that his definition *defines* causal effect. This would mean that causal effects are what epidemiologists tell us they are in a sort of metaphysically unsatisfying and somewhat patronizing way.

Let me refer briefly to an exchange from the epidemiological literature about the capability of ABMs to contribute to causal inference. Marshall and Galea have argued that ABMs “represent a promising novel approach to identify and evaluate complex causal effects” (Marshall and Galea 2015). Although they refer to causal inference in this quote and in the title of their paper, the authors seem to avoid this notion in the body of the paper and refer instead to the exploration, elucidation, and interrogation of the causal relationships modeled in an ABM. Their argument rests on the capability of ABMs to represent multiple causal interrelations (their view of a complex system):

We argue that agent-based modeling offers an alternative and complementary approach to elucidate complex causal interdependencies that are of interest in epidemiology. Specifically, the forms of the relationships among causes (which are broadly defined here and can include agent traits as well as environments) are operationalized by the rules Z . The rule set consisting of functions $f()$, $g()$, and $h()$ can include nonlinear components, including feedback loops, such that linear independence need not be assumed. By altering the rule set Z and running the simulation under different assumed causal relationships and processes, the effect(s) of interdependent (i.e., joint) exposures can be explored and interrogated (Marshall and Galea 2015: 96).

Marshall and Galea call the causal interrelationships they are interested in *complex*. I take it as implicit that by this they refer to complex systems, not just complicated ones. They stress the possibility to model non-linear relationships—a characteristic of complex systems. Thus, their view seems to be in keeping with the notion discussed in §5 below that the sheer complexity of interactions of agents in ABMs may give rise to emergent phenomena. More importantly, it is the intervention by the modeler (altering the rule set under different causal assumptions) that renders the ABM a helpful tool in causal exploration and interrogation, to use Marshall and Galea’s terms. This view grants epistemological value to ABMs based on the possibility to *manipulate* them and explore the consequences, which resonates with interventionist accounts of causation.

One invited commentator, Ana Diez-Roux, disagrees with the notion that ABMs can help with causal inference in epidemiology (Diez Roux 2014). The following excerpt from her abstract puts her position, which I see as one point of departure for my proposal in §6 below, in a nutshell:

As discussed by Marshall and Galea [...], systems approaches are appealing because they allow explicit recognition of feedback, interference, adaptation over time, and nonlinearities. However, they differ fundamentally from the traditional approaches to causal inference used in epidemiology in that they involve creation of a virtual world. Systems modeling can help us understand the plausible

implications of the knowledge that we have and how pieces can act together in ways that we might not have predicted. [...] However, the validity of any causal conclusions derived from systems models hinges on the extent to which the models represent the fundamental dynamics relevant to the process in the real world. For this reason, systems modeling will never replace causal inference based on empirical observation. Causal inference based on empirical observation and simulation modeling serve interrelated but different purposes (Diez-Roux 2014: 100).

Of note, Diez-Roux does not say that ABMs are incapable of helping with causal inference in principle. She only says that ABM-generated models are not like epidemiological approaches to causal inference based on observed data. However, I agree with her notion that simulated data from ABMs are epistemologically inferior to observational epidemiological data simply because the underlying data are not real-world data but data generated *in silico*.

4. Interaction and Emergence

Let us now move on to the question whether the system behavior of ABMs can be reduced to the interactions among agents' characteristics and behaviors or if it is an *emergent* phenomenon. The question I am interested in is about the relationship between mechanistic explanation and emergence. In brief, if ABMs are a non-deterministic black-box and the system behavior they exhibit is truly emergent, what remains of the notion that ABMs represent causal-mechanistic explanations? What kind of causal mechanism would be explained by an ABMs whose inner workings remain in the dark and whose results are *by definition* unpredictable and surprising? (I see a similarity here to the current discussion about the transparency, explainability, and interpretability of machine learning algorithms (Roscher et al. 2020), but an exhibit of this parallel will have to wait for another day.) On the other hand, if ABMs really provide causal-mechanistic explanations we should be able to predict the phenomena they generate, which would render them non-emergent.

4.1 Emergence Defined

The classic reference on emergence, published by Jeffrey Goldstein in the first issue of the journal of the same name, defines emergence as

the arising of novel and coherent structures, patterns, and properties during the process of self-organization in complex systems. Emergent phenomena are conceptualized as occurring on the macro level, in contrast to the micro-level components and processes out of which they arise (Goldstein 1999: 49).

Think of a complex system as having a micro level (components) and a macro level (surface). Goldstein defines emergent phenomena as (1) radically novel, (2) coherent, (3) macro-level, (4) dynamic, and (5) ostensive. *Radical novelty* refers to the fact that emergent phenomena appear at the macro level without having previously been present in the complex system under study and cannot be derived from or predicted based on knowledge about what is going on at the micro-level. *Coherence* means that emergent properties maintain "some sense of identity over time" (*ibid.*), *macro-level* means that emergence is observable at the surface level of the observed system, not the micro-level constituted by its components, *dynamic*

refers to emergent phenomena as not preformed but as developing over time, and *ostensive* as being recognized by “showing themselves”.

Most important for our present discussion, however, is that Goldstein sees one of the main roles of emergence in science as explanatory:

In respect to its use in scientific explanation, the construct of emergence is appealed to when the dynamics of a system seem better understood by focusing on across-system organization rather than on the parts or properties of parts alone (Goldstein 1999: 50).

Thus, in keeping with Goldstein’s characterization of emergent phenomena, although their occurrence on the macro-level is *produced* by what is going on at the micro-level, they come “out of the blue” because they do not *depend* on the behavior of individual micro-level variables (agents in ABMs) but on the overarching function of the whole system. Thus, if ABMs are truly complex (non-deterministic, non-linear) systems, they would produce emergent effects at the output level that are not predicted, or even *predictable*, by means of applying knowledge about the agents and their interactions. In contrast, these results would be ostensive occurrences that rely on the function of all interacting parts. The point here is that ABMs yield models of mechanisms that do not necessarily represent any real-world mechanism, be it biological or social mechanisms. It represents only itself, based on input conditions and probabilistic rules for agent interactions and status changes. If an ABM yields an outcome, be it emergent or expected, the occurrence of that outcome can then be explained by analyzing the workings of the modeled mechanism in silico.

What kind of mechanism consists of interactions between parts over time but is not “productive of regular changes” (per Machamer et al.’s definition) but instead to radically novel, dynamic, and ostensive phenomena? Can ABMs explain mechanisms or emergence, or both?

4.2 Weisberg: Mechanistic Explanations vs Emergence Explanations

The question whether ABMs can explain emergent phenomena is what Weisberg considers “the most controversial claim about IBMs [...] Not everyone is convinced” (Weisberg 2014: 788). He quotes ecologist Joan Roughgarden as saying that she doesn’t “think it’s easy to discern the causation being revealed by an IBM simulation. And if we don’t learn something about causation we don’t learn anything scientifically important” (personal communication quoted in Weisberg 2014). (Of note, Weisberg and Roughgarden’s IBMs and our ABMs are the same thing; see above.)

Weisberg suggests a distinction between explanations of emergent phenomena (mechanistic explanations) and explanations of the emergence of phenomena (emergence explanations). On his view, mechanistic explanations provide a “generalized mechanistic understanding of the dependence of higher-level properties and patterns on lower-level mechanistic factors” (Weisberg 2014: 789). I take this to mean an explanation that is based on the description of the elements of a mechanism and their interactions as being what *somehow* leads to an emergent phenomenon. He shows how certain causal graphs (relational depictions of phenomena in boxes with causal arrows between them) can depict the relationships among micro-level factors that can help generate mechanistic explanations. Interestingly,

the *kind* of causal graph he chooses suggests that on his view ABMs can model *biological* mechanisms because the causal mechanism depicted in his example permits feedback loops, an important characteristic of mechanisms in biological explanations (Bechtel 2011). In contrast, the directed acyclic graphs (DAGs) that are frequently used in epidemiological causal reasoning do *not* allow feedback loops, a feature preferred in causal reasoning because the vertices can be ordered, simplifying causal argumentation immensely. No such topological order is possible in cyclic graphs (Dasgupta, Papadimitriou, and Vazirani 2008: 96).

Emergence explanations, on the other hand, would require us to provide “reductive explanations that show how emergent phenomena arise from lower-level interactions” (Weisberg 2014: 792). They would require us to clarify the *somehow* that generates an emergent phenomenon. But one main problem with both cyclic and acyclic graphs is that it is unclear *what exactly the arrows represent*. If it is true that causation is “one word, many things” and that “there are different kinds of causal relations imbedded in different kinds of systems” (Cartwright 2004: 805), the edges (arrows) between different vertices (characteristics of agents in ABMs) would potentially represent different sub-mechanisms. I read Weisberg as saying that we cannot use ABMs to provide emergence explanations unless we can specify exactly what is in each of these arrows, and I agree with him on that. On the other hand, he seems to say that ABMs can provide mechanistic explanations. Let me add that if all mechanisms are causal, I assume that Weisberg would conclude that ABMs can provide causal-mechanical explanations and I would agree with him on that as well.

I also suggest that his usage of non-DAGs to depict what ABMs model not only fits biological but also social mechanisms. Note that the Columbia account of ABMs above explicitly mentions feedback loops. Indeed, some research on COVID-19 has revealed interesting feedback loops even across scales of representation (micro-level, macro-level). For example, one computational study suggests that macro-level dynamics such as social distancing can result in micro-level changes all the way down in the genetic makeup of SARS-CoV-2 (Barrett et al. 2021).

But perhaps, at least in the context of ABMs, we shouldn’t ask too much of the arrow semantics in causal graphs, for in ABMs the relationship between all agents and all their characteristics is simply a mathematical relationship, not a biological one. This brings us to the next notion reflected in the Columbia account of ABM, impossible interactions.

5. Impossible Interactions

A major motivation to use ABMs comes from their flexibility to be manipulated in ways no observational or interventional epidemiological study could be manipulated. In essence, ABMs can be used to model the “impossible” because the characteristics of agents are variables created *for* the model and *by* the model. Furthermore, there is only one kind of relationship between and among variables in ABMs, a mathematical relationship represented by stochastic functions.

Based on the findings in the systems biology and population health/sociological/ecological literature discussed in the previous sections we can postulate that ABMs are considered models of social mechanisms. Such mechanisms are modeled in ABMs by creating *interactions* between agent’s characteristics among

each other and between agents' and their environment's characteristics. How does this look like *inside* an ABM?

5.1 Interactions

The term *interaction* is most often used in ABMs to denote the narrowing of virtual physical space between two agents to a level at which a status change occurs in at least one of them (Winkelmann et al. 2021). Based on certain parameters, each individual agent will move through virtual space until a pre-programmed fit between a set of characteristics of two agents leads to contact and infection with a certain prespecified likelihood. At this point the status of the heretofore "uninfected" agent switches to "infected". Because such status changes are dependent on certain constellations of variables at certain timepoints, and because these constellations are derived from a whole set of characteristics assigned to agents in a stochastic fashion, these interactions and associated status changes are *not* predetermined. In this sense, ABMs are non-deterministic, and each run of the model will yield a slightly different end result. Many runs need to be performed to narrow down the probability distribution of results at the macro-level. At the population level, population wide parameters such as "infection prevalence" change from starting conditions to a different value over the duration of model run time, depending on how many individuals will be newly infected (incidence) while the model is running. Such result is sometimes considered "emergent" since it is not fully determined by model parameters in an equation-like fashion.

In the above scenario, the interaction is between two agents. Interactions can also occur between agents and the spatial environment. For example, certain areas in the virtual space can be designated as different in terms of social characteristics (e.g., high crime, low crime, no crime regions) and the likelihood of a status change of an agent (e.g., becoming the victim in a street mugging) would be different in these different regions. Moreover, agent-agent interactions could be modeled as representing just such a mugging (or not) and differ by section of the virtual space.

5.2 The Impossible

These considerations highlight one of the oft-praised advantages of ABMs, the possibility to design interactions *in any way* the modeler desires, even impossible or unethical ones. In essence, the functions of ABMs are completely devoid of the need for plausibility and ethical considerations. Nothing prevents the design of an ABM of a randomized controlled trial of the effect of COVID-19 infection on survival. Obviously, although such trial would be possible in principle, it would (luckily) never be approved by an institutional review board.

But aside from being a potential tool for modeling the unethical, another important possibility is to model mechanistic relationships across levels along the bio-psycho-social spectrum. Agent-based models can evaluate interactions among and between agents and their environments regardless of a known mechanism between, say, agents' socioeconomic background, their immune status, and their risk of SARS-CoV-2 infection. The flip side of ABMs' *inability* to provide Weisbergian emergence explanations is the benefit for the modeler to simply ignore the *somehow* expected from such explanations without sacrificing the capability of their model to provide causal-mechanical explanations.

5.3 ABMs as Multiscale Models

In epidemiological research, multilevel modeling that integrates variables across the individual, household, and community level is a common approach. Such models are called *multi-scale* or *nested* models and have become common in infectious disease modeling (Hart et al. 2020). Multiscale models have traditionally been based on integro-differential equations (IDEs), but the usage of ABMs has recently become more frequent. Such models can easily integrate the interaction between biological and behavioral processes at the level of the individual and social processes at the population level.

At least some philosophers seem to feel comfortable with the idea of trans-level interaction and state that “our health is not just a metabolic response to toxins; it is about a complex social and biological interaction—a relational process or mechanism” (Parkkinen et al. 2018). Indeed, I suggest that agent-based multiscale models can provide the proposed integration of biological, behavioral, and social mechanism in a concept that Kelly, Kelly, and Russo have advocated for and called *mixed mechanisms* (Kelly, Kelly, and Russo 2014). However, I think that they can do even more: they can explore comprehensive etio-prognostic explanations of illness occurrence, development, and prognosis. Indeed, ABMs can simulate not only the joint activities of determinants of illness occurrence (causes and mechanisms) in etiological explanations (Dammann 2020), but also the joint activities of the determinants of the clinical course (disease development) and its outcomes (cure, death, or anything in between). They can even include the potential impact of etiological contributors such as *conditions* that are different from causes in non-trivial ways (Broadbent 2008) that I regret not being able to rehearse here in detail. In the next and final section, I propose that while ABMs’ role in causal inference might be limited, they can provide *etio-prognostic explanations* by integrating determinants of illness occurrence (etiology) as well as determinants of disease development and outcome (prognosis).

6. ABMs as Etio-Prognostic Explanations

Above, I have rejected the idea that ABMs can help with causal inference, but support the notion that ABMs can be helpful as explanations of causal-mechanical (etiological) processes of illness occurrence. Moreover, I propose that they can help even further by simulating the trajectory of illness development and outcome. Let me begin by outlining *etiological explanations* (Dammann 2017, 2020) and what I mean by *etio-prognostic explanation*.

6.1 Etiological Explanations

In epidemiology, an obsession with causal inference abounds. The main idea seems to be that epidemiological methods can provide an apparatus that allows for causal inference based on observational epidemiological data. The underlying assumptions appear to be that observed statistical associations are not to be considered reflective of a causal relationship unless they come from ideal randomized experiments (Hernán 2004). A simple and straight forward rejection of this proposal would need to show that ideal randomized experiments do not exist. Indeed, some philosophers have offered this argument as well as other considerations that should reduce our confidence in causal inference from randomized clinical trials, the gold standard of the randomized experiment in clinical epidemiology (Worrall 2007,

Cartwright 2007, 2010, Deaton and Cartwright 2016). If these arguments, which I cannot fully discuss here for reasons of space, carry any weight, there may just not be any way to reliably infer causality from epidemiological data. Instead of making causal inference the holy grail of epidemiological research, a gentler, less exclusive perspective can be taken according to which epidemiology contributes to the generation of etiological explanations, which refer to purported causes of illness, the mechanisms they initiate, and the disease (illness) that occurs. This theoretical model of illness occurrence is a process model, with causation process and disease process overlapping and jointly representing the etiological process (Figure 1). Providing such etiological explanation means providing a coherent set of hypotheses that support the observed data, explaining the occurrence of the disease and its clinical outcome (for a philosophical take on explanatory coherence in epidemiology, see Dammann 2018).

Comprehensive etiological explanations may include reference to initiators (causes), mediators, modifiers (both part of the pathogenetic mechanism), and facilitators. Causes (e.g., Sars-CoV-2 infection) are factors that initiate the mediating pathomechanism (e.g., severe inflammation in the lung) which leads to pulmonary disease, sometimes respiratory failure, and death (outcome). Modifiers in this explanation are factors that change the impact of causes and mechanisms (e.g., vaccination or social distancing), while facilitators are any biological, behavioral, or societal conditions that have an impact on the remainder of the etiological process (such as age, race, access to healthcare, and so forth). Modeling such comprehensive etiological explanation is exactly what I see multi-scale ABMs as capable of doing. They can simulate what might happen in a population given a certain constellation of characteristics that describe the interactions between initiators/causes, mediators/mechanisms, modifiers of the causation process, and facilitators/background conditions.

6.2 Etio-Prognostic Explanations

Etiological explanations are explanations that tell a cogent story of illness occurrence that is justified by reference to coherent causal and mechanistic evidence. Giving an etiological explanation means to provide a list of causes (even if the list has only one item) and mechanisms that, taken together, suffice to change the beliefs of the hitherto unconvinced about why and how the illness occurred. I think that this characterization of etiological explanations should work in both medical (single patient) settings as well as in epidemiological (population) contexts. Agent-based models that provide etiological explanations would be models of the entire etiological process from cause via mechanism to clinical disease as depicted in Figure 1. Any ABM that models COVID-19 infection incidence would provide an etiological explanation.

However, many ABMs that have been developed to model population-wide aspects of the pandemic do more: they also include estimates of hospitalizations based on estimates of illness severity, admission to intensive care, and mortality, as in the example provided above. These kinds of ABM not just explain illness occurrence but also what happens afterwards, the *prognosis* of illness. Let me offer the following table to make some potentially helpful distinctions.

Explanation →	Causal	Mechanical	Clinical	Prognostic
Explanans	Causes (risk factors)	Pathogenesis (biology)	Clinical course (signs and symptoms)	Outcome (cure, death, or anything in between)
Explanandum	Why (“roots”)	How?	Clinical presentation	Prognosis
Source of evidence	Epidemiology	Biosciences	Clinical medicine	Follow up (medicine, epidemiology)
	Etiological Explanation			
			Prognostic Explanation	
	Etio-Prognostic Explanation			

Table 1. Characteristics of causal, mechanical, clinical, and prognostic explanations.

Of note, the “intended *explicandum* [of scientific explanations] is, very roughly, explanations of *why* things happen, where the ‘things’ in question can be either particular events or something more general—e.g., regularities or repeatable patterns in nature” (Woodward and Ross 2021). I am aware that explaining *why* something happens is a very different thing than explaining its *consequences*. Indeed, such an explanation would probably not be considered *scientific*. However, a slight change of perspective might allow us to reintroduce science through the backdoor. We could say that what happens after the initial occurrence of illness is just the *occurrence* of aspects of disease development and outcome. Thus, the prognostic part of etio-prognostic explanations can be viewed as providing a plain old etiological explanation. This way, one could see prognostic explanations as scientific, i.e., by recognizing them as etiological explanations of a different target entity.

However, I am interested in the mere *practical* usefulness of explanations of illness occurrence and outcome. I prefer looking at ABMs as providing a pragmatic kind of explanation, which is simply helpful by illuminating both the *etiology and prognosis* of illness. This is exactly what we expect from ABMs in the context of the COVID-pandemic: explanations why and how illness occurrence patterns arise at the population level, how they evolve, and what their consequences are.

7. Conclusion

In this paper, I have discussed the epistemological characteristics of ABMs, one type of simulation model used in the context of the COVID-19 pandemic. In contrast to equation-based models, ABMs are algorithms that use individual agents and attribute changing characteristics to each one, multiple times during multiple iterations over time. Based on my discussion, I conclude that ABMs can explain causal mechanisms but cannot provide emergence explanations, because they cannot provide information about exactly why low-level phenomena give rise to those emergent phenomena. This is also one reason why I believe that ABMs cannot help with causal inference. Another reason is that ABMs do not reflect

real-world processes but the causal-mechanical intuitions of the modeler. On the other hand, ABMs can integrate “impossible” multi-scale interactions between initiators, mediators, moderators, and conditions, and may be useful as comprehensive etio-prognostic explanations of illness occurrence and outcome.

References

- Aleta, A., Martin-Corral, D., Piontti, Y., Pastore, A., Ajelli, M., Litvinova, M., Chinazzi, M., Dean, N.E., Halloran, M.E., Longini, I.M. Jr., Merler, S., Pentland, A., Vespignani, A., Moro, E., and Moreno, Y. 2020, “Modelling the Impact of Testing, Contact Tracing and Household Quarantine on Second Waves of COVID-19”, *Nature Human Behaviour*, 4, 9, 964-71, doi: 10.1038/s41562-020-0931-9.
- Barrett, C., Bura, A.C., He, Q., Huang, F.W., Li, T.J.X., Waterman, M.S., and Reidys, C.M. 2021, “Multiscale Feedback Loops in SARS-CoV-2 Viral Evolution”, *Journal of Computational Biology*, 28, 3, 248-56. doi: 10.1089/cmb.2020.0343.
- Bechtel, W. 2011, “Mechanism and Biological Explanation”, *Philosophy of Science*, 78, 533-57.
- Broadbent, A. 2008, “The Difference between Cause and Condition”, *Proceedings of the Aristotelian Society*, 108, 1, pt 3, 355-64.
- Cartwright, N. 1989, *Nature’s Capacities and Their Measurement*, Oxford-New York: Clarendon Press.
- Cartwright, N. 2004, “Causation: One Word, Many Things”, *Philosophy of Science*, 71, 805-19.
- Cartwright, N. 2007, “Are RCTs the Gold Standard?”, *Biosocieties*, 1, 11-20.
- Cartwright, N. 2010, “What Are Randomised Controlled Trials Good For?”, *Philosophical Studies*, 147, 1, 59-70.
- Casini, L., and Manzo, G. 2016, “Agent-Based Models and Causality: A Methodological Appraisal”, *The IAS Working Paper Series*, 7, 1-80.
- Conte, R., and Paolucci, M. 2014, “On Agent-Based Modeling and Computational Social Science”, *Frontiers in Psychology*, 5, doi: 10.3389/fpsyg.2014.00668.
- Cuevas, E. 2020, “An Agent-Based Model to Evaluate the COVID-19 Transmission Risks in Facilities”, *Computers in Biology and Medicine*, 121, doi: 10.1016/j.compbiomed.2020.103827.
- Dammann, O. 2017, “The Etiological Stance: Explaining Illness Occurrence”, *Perspectives in Biology and Medicine*, 60, 2, 151-65, doi: 10.1353/pbm.2017.0025.
- Dammann, O. 2018, “Hill’s Heuristics and Explanatory Coherentism in Epidemiology”, *American Journal of Epidemiology*, 187, 1, 1-6, doi: 10.1093/aje/kwx216.
- Dammann, O. 2020, *Etiological Explanations: Illness Causation Theory*, Boca Raton: CRC Press.
- Dasgupta, S., Papadimitriou, C.H., and Vazirani, U.V. 2008, *Algorithms*, Boston: McGraw-Hill Higher Education.
- Deaton, A., and Cartwright, N. 2016, “Understanding and Misunderstanding Randomized Controlled Trials”, *National Bureau of Economic Research Working Paper Series*, No. 22595, doi: 10.3386/w22595.

- Diez Roux, A.V. 2014, “Invited Commentary: The Virtual Epidemiologist—Promise and Peril”, *American Journal of Epidemiology*, 181, 2, 100-102, doi: 10.1093/aje/kwu270.
- Goldstein, J. 1999, “Emergence as a Construct: History and Issues”, *Emergence*, 1, 1, 49-72, doi: 10.1207/s15327000em0101_4.
- Hart, W.S., Maini, P.K., Yates, C.A., and Thompson, R.N. 2020, “A Theoretical Framework for Transitioning from Patient-Level to Population-Scale Epidemiological Dynamics: Influenza A as a Case Study”, *Journal of The Royal Society Interface*, 17, 166, doi: 10.1098/rsif.2020.0230.
- Hedström, P. and Swedberg, R. 1998, *Social Mechanisms: An Analytical Approach to Social Theory*, in Elster, J. and Hernes, G. (eds.), *Studies in Rationality and Social Change*, Cambridge-New York: Cambridge University Press.
- Hernán, M.A. 2004, “A Definition of Causal Effect for Epidemiological Research”, *Journal of Epidemiology and Community Health*, 58, 4, 265-71.
- Hoertel, N., Blachier, M., Blanco, C., Olfson, M., Massetti, M., Sánchez Rico, M., Limosin, F., and Leleu, H. 2020, “A Stochastic Agent-Based Model of the SARS-CoV-2 Epidemic in France”, *Nature Medicine*, 26, 9, 1417-21, doi: 10.1038/s41591-020-1001-6.
- Hunter, E., Mac Namee, B., and Kelleher, J.D. 2017, “A Taxonomy for Agent-Based Models in Human Infectious Disease Epidemiology”, *Journal of Artificial Societies and Social Simulation*, 20, 3, doi: 10.18564/jasss.3414.
- Kelly, M. P., Kelly, R.S., and Russo, F. 2014, “The Integration of Social, Behavioral, and Biological Mechanisms in Models of Pathogenesis”, *Perspectives in Biology and Medicine*, 57, 3, 308-28, doi: 10.1353/pbm.2014.0026.
- Kerr, C.C., Robyn, M., Stuart, D.M., Romesh, G.A., Rosenfeld, K., Hart, G.R., Núñez, R.C., Cohen, J.A., Selvaraj, P., Hagedorn, B., George, L., Jastrzębski, M., Izzo, A., Fowler, G., Palmer, A., Delpont, D., Scott, N., Kelly, S., Bennette, C.S., Wagner, B., Chang, S., Oron, A.P., Wenger, E., Panovska-Griffiths, J., Famulare, M., and Klein, D.J. 2021, “Covasim: An Agent-Based Model of COVID-19 Dynamics and Interventions”, *medRxiv*, <https://www.medrxiv.org/content/10.1101/2020.05.10.20097469v3>
- Machamer, P.K., Darden, L. and Craver, C.F. 2000, “Thinking About Mechanisms”, *Philosophy of Science*, 67, 1-25.
- Marshall, B.D., and Galea, S. 2015, “Formalizing the Role of Agent-Based Modeling in Causal Inference and Epidemiology”, *American Journal of Epidemiology*, 181, 2, 92-99, doi: 10.1093/aje/kwu274.
- Medicine, Institute of 2003, *The Future of the Public’s Health in the 21st Century*.
- Parkkinen, V., Wallmann, C., Wilde, M., Clarke, B., Illari, P., Kelly, M.P., Norell, C., Russo, F., Shaw, B., and Williamson, J. 2018, *Evaluating Evidence of Mechanisms in Medicine*, Cham: Springer.
- Roscher, R., Bohn, B., Duarte, M.F., and Garcke, J. 2020, “Explainable Machine Learning for Scientific Insights and Discoveries”, *IEEE Access*, 8, 42200-42216, doi: 10.1109/access.2020.2976199.
- Shamil, M.S., Farheen, F., Ibtehad, N., Khan, I.M., and Sohel Rahman, M. 2021, “An Agent-Based Modeling of COVID-19: Validation, Analysis, and Recommendations”, *Cognitive Computation*, doi: 10.1007/s12559-020-09801-w.
- Silva, P.C.L., Batista, P.V.C., Lima, H.S., Alves, M.A., Guimarães, F.G., and Silva, R.C.P. 2020, “COVID-ABS: An Agent-Based Model of COVID-19 Epidemic to

- Simulate Health and Economic Effects of Social Distancing Interventions”, *Chaos, Solitons & Fractals* 139, doi: 10.1016/j.chaos.2020.110088.
- Staffini, A., Svensson, A.K., Chung, U., and Svensson, T. 2021, “An Agent-Based Model of the Local Spread of SARS-CoV-2: Modeling Study”, *JMIR Medical Informatics*, 9, 4, doi: 10.2196/24192.
- Stinchcombe, A.L. 1991, “The Conditions of Fruitfulness of Theorizing About Mechanisms in Social Science”, *Philosophy of the Social Sciences*, 21, 3, 367-88, doi: 10.1177/004839319102100305.
- Susser, M. 1991, “What is a Cause and How Do We Know One? A Grammar for Pragmatic Epidemiology”, *Am J Epidemiol*, 133, 635-48.
- Topping, C.J., Høye, T.T., and Olesen, C.R. 2010, “Opening the Black Box—Development, Testing and Documentation of a Mechanistically Rich Agent-Based Model”, *Ecological Modelling*, 221, 2, 245-55, doi: 10.1016/j.ecolmodel.2009.09.014.
- Tracy, M., Cerdá, M., and Keyes, K.M. 2018, “Agent-Based Modeling in Public Health: Current Applications and Future Directions”, *Annual Review of Public Health*, 39, 1, 77-94, doi: 10.1146/annurev-publhealth-040617-014317.
- Truszkowska, A., Behring, B., Hasanyan, J., Zino, L., Butail, S., Caroppo, E., Jiang, Z., Rizzo, A., and Porfiri, M. 2021, “High-Resolution Agent-Based Modeling of COVID-19 Spreading in a Small Town”, *Advanced Theory and Simulations*, 4, 3, doi: 10.1002/adts.202000277.
- Van Dyke Parunak, H., Savit, R., and Riolo, R.L. 1998, *Agent-Based Modeling vs. Equation-Based Modeling: A Case Study and Users’ Guide*, in Sichman, J.S., Conte, R., and Gilbert, N. (eds.), *Multi-Agent Systems and Agent-Based Simulation*, Berlin: Springer, 10-25.
- Venkatramanan, S., Lewis, B., Chen, J., Higdon, D., Vullikanti, A., and Marathe, M. 2018, “Using Data-Driven Agent-Based Models for Forecasting Emerging Infectious Diseases”, *Epidemics*, 22, 43-49, doi: 10.1016/j.epidem.2017.02.010.
- Weisberg, M. 2014, “Understanding the Emergence of Population Behavior in Individual-Based Models”, *Philosophy of Science*, 81, 5, 785-97, doi: 10.1086/677405.
- Winkelmann, S., Zonker, J., Schütte, C., and Conrad, N.D. 2021, “Mathematical Modeling of Spatio-Temporal Population Dynamics and Application to Epidemic Spreading”, *Mathematical Biosciences*, 336, doi: 10.1016/j.mbs.2021.108619.
- Wolkenhauer, O. 2014, “Why Model?”, *Frontiers in Physiology*, 5, 21, doi: 10.3389/fphys.2014.00021.
- Woodward, J. and Ross, L. 2021, “Scientific Explanation”, in *The Stanford Encyclopedia of Philosophy*, Zalta, E.N. (ed.).
- Worrall, J. 2007, “Why There’s no Cause to Randomize”, *British Journal for the Philosophy of Science*, 58, 3, 451-88, doi: 10.1093/Bjps/Axm024.