



What are general models about?

Alkistis Elliott-Graves¹ 

Received: 30 January 2021 / Accepted: 26 November 2022 / Published online: 13 December 2022

© The Author(s) 2022

Abstract

Models provide scientists with knowledge about target systems. An important group of models are those that are called *general*. However, what exactly is meant by generality in this context is somewhat unclear. The aim of this paper is to draw out a distinction between two notions of generality that has implications for scientific practice. Some models are general in the sense that they apply to *many systems in the world* and have *many particular targets*. Another sense is captured by models that are aimed at understanding the fundamental or underlying dynamics of a phenomenon, as opposed to how it manifests in each particular case. They have non-specific, i.e. *generic targets*. While both notions of generality and genericness are legitimate and correspond to different aspects of scientific practice, they must be distinguished. Failing to do so obscures the danger of overgeneralisation faced by general models and facilitates the illegitimate use of generic models as general models. This can lead to a reduction of the explanatory and predictive power of both.

Keywords Generality · Scientific models · Target system · Ecological modelling · Segregation

1 Introduction

A central question in modelling epistemology is: “in virtue of what do models provide us with knowledge about the world?”. This question is particularly interesting in cases where there is not a one-to-one mapping between a particular real-world system and a specific model designed to provide information about that system. One such case is that of *general models*. There are numerous conceptions of ‘generality’ with respect to models (Cooper, 1998; Evans et al., 2013; Kuorikoski & Lehtinen, 2009; Levy, 2018; Matthewson & Weisberg, 2009; Weisberg, 2013; Ylikoski & Aydinonat, 2014). In this paper, I will focus on two of these conceptions, *generality* and *genericness*, and argue that they should be kept distinct. At the conceptual level,

✉ Alkistis Elliott-Graves
a.elliott-graves@uni-bielefeld.de

¹ Abteilung Philosophie, Universität Bielefeld, Postfach 100 131, 33501 Bielefeld, Germany

this distinction helps us to understand the above question, as it sheds light on what general models are *about*. Yet this discussion also has practical implications, as it can help to elucidate two common problems in the practice of modelling.

There is a consensus in the philosophical literature that models can be used for different purposes, including but not limited to explanation, prediction and intervention (Giere, 2004; Knuuttila & Loettgers, 2016a, 2016b; Morgan & Morisson, 1999; Weisberg, 2013). Scientists and philosophers often include a model's intended use in their discussions of model-world relations, model success and a model's source of value. For example, the logistic growth model in Ecology is used to explain and predict the size of many actual populations of plants and animals (Barlow, 1992; Brashares et al., 2010; Levins, 1966). One of its sources of value, perhaps the most important one, is that applies to many diverse actual systems in the world (Levins, 1966; Orzack & Sober, 1993). When discussing Schelling's model of residential segregation, however, philosophers and social scientists identify a different use and source of value. This model is used to acquire an understanding of the fundamental or underlying dynamics of a phenomenon (segregation), as opposed to how it manifests in each particular case (Aydinonat, 2007; Kuorikoski & Lehtinen, 2009; Weisberg, 2013; Ylikoski & Aydinonat, 2014). Moreover, its source of value is not that it provides information about how segregation actually comes about in real cities, but because it provides a *possible* explanation of how segregation *could* come about (Aydinonat, 2007; Kuorikoski & Lehtinen, 2009; Ylikoski & Aydinonat, 2014).

Despite identifying these different uses and sources of value for models, philosophers of science use the term 'general' for both types of model use, routinely grouping both types together under a single characterisation, usually 'general models', sometimes also 'minimal models' or 'toy models' (Nguyen, 2020; Reutlinger et al., 2018; Weisberg, 2013). I believe that this undermines the importance of distinguishing between the two types of model use. In what follows, I will argue that we should instead highlight the differences between these two types of model use, and that this can be achieved by distinguishing between the concepts of *generality* and *genericness*.

I should point out that my argument pertains to model *use* rather than a model's *nature*, i.e. what models are about rather than what they are. I will not be arguing that a model itself is general or generic but will investigate what happens when it is used generally or generically in certain contexts. When, on occasion, I refer to a general or generic model, this is just shorthand for a model being used generally or generically. It may seem that framing the argument in terms of model use undermines the distinction between generality and genericness, but this is not the case (see Sect. 3.2). Even if a model can be used generally and/or generically, this does not mean that it should. The point of the distinction is to provide the conceptual foundation for determining which model use is best suited to certain contexts and to explain why one use could be ill-advised in a certain context (see Sect. 5).

Another common view in the literature is that models are about *target systems* (Frigg, 2009; Giere, 2004; Knuuttila, 2005; Matthewson & Weisberg, 2009; Suárez, 2010; Weisberg, 2013). Target systems can also be understood in a number of different ways, e.g. as parts of real-world systems, abstract systems, or even hypothetical systems (Elliott-Graves, 2020; Suárez, 2010; Weisberg, 2013). However, when

discussing general models, philosophers view their targets as ‘generalised abstract systems’ (Aydinon, 2007; Kuorikoski & Lehtinen, 2009; Weisberg, 2013; Ylikoski & Aydinon, 2014). I will argue that the distinction between generality and genericness extends to target systems, while the concept of a ‘generalised abstract system’ does not fully capture either type of target. Furthermore, general models apply to many particular real-world targets, whereas generic models apply to generic targets.

I start by examining a general model (the logistic growth model) and a generic model (Schelling’s segregation model) (Sect. 2). I further clarify the distinction between generality and genericness (Sect. 3), focusing on the concept of *abstraction* (3.1), the importance of model *use* and the possibility that a single model can be used both generally and generically (3.2), the case of model families or types (3.3), models’ source of value (3.4) and the distinction between *a* and *p* generality (3.5). I then examine the target systems of general and generic models, arguing that they are also distinct (Sect. 4). I show, contra Weisberg (2013), that neither general nor generic models have ‘generalised target systems’. Instead, general models have many local targets while the targets of generic models are themselves generic. In the final Sect. (5), I argue that two of the most important dangers in modelling, (i) overgeneralisation and the (ii) illegitimate transformation of how-possibly to how-actually explanations, are facilitated by the failure to distinguish between general and generic model use.

2 Two ‘general’ models

2.1 Logistic growth

One of the most significant models in the history of Ecology is the *logistic growth model*, invented independently by Pierre Verhulst and Raymond Pearl, to describe how populations grow (Kingsland, 1982). The model measures how the growth rate of a population (N) is limited by the density of the population itself (the differential form of the model description can be seen in Eq. 1, below):

$$\frac{dN}{dT} = rN \left(1 - \frac{N}{K} \right) \quad (1)$$

Here, (r) is the intrinsic growth rate, the maximum possible growth rate of the population. It is roughly equivalent to the number of deaths in the population subtracted from the number of births in that population. (K) the *carrying capacity* of the environment, refers to the maximum number of individual organisms that a particular environment can support. The most important factors that affect (K) are resources, which can vary across environments and in terms of how they are used by different species. Populations displaying logistic growth have an initial phase of rapid growth, where the growth rate increases exponentially. As resources are used up, the growth rate slows, until the carrying capacity is reached, whereupon the growth rate stabilises.

From its inception, the logistic growth model was intended to be general –it was meant to apply to many different populations. Box 1 illustrates how Reed and Pearl demonstrated that the shape of the uninstantiated version of the model was almost identical to four different instantiations, including three different human populations (New York, New Jersey and Connecticut), and a population of drosophila in a laboratory. The model was also applied to many other human populations, including those of Germany, Denmark and Japan (Kingsland, 1982).

Generality as ‘applicability to many systems’ is the standard notion of generality in Ecology (Evans et al., 2013) and is also quite common in the philosophical accounts of scientific modelling, especially in the literature responding to and influenced by the work of Richard Levins (Cooper, 1998; Levins, 1966, 1993; Matthewson & Weisberg, 2009; Odenbaugh, 2003; Orzack & Sober, 1993; Weisberg, 2013).¹ Levins argued that generality is one of three desiderata for ecological models, the other two being realism and precision (Levins, 1966, 1993). Generality, in this sense, is valuable because it helps scientists make sense of seemingly disparate phenomena or systems. Showing that a particular phenomenon is an instance of a more general pattern because it is caused by a common set of rules or mechanisms can help scientists explain the phenomenon or predict how the system will behave in the future. This, in turn, allows scientists to transfer the knowledge they gain from investigating one system to the others, to use the model to predict the future behaviour of these systems and have quite high confidence that our predictions will be accurate. In fact, many accounts of scientific explanation have highlighted generalisability through theoretical unity as the hallmark of good scientific explanations (Hempel & Oppenheim, 1948; Kitcher, 1981). In the case of logistic growth, applying the model to many different populations was possible because its fundamental dynamic (exponential growth curbed by limiting resources) was thought to be the core factor of population growth, present in all real-world populations. In other words, generality helps to uncover the relevant similarities between seemingly disparate systems. In fact, the early success of the logistic growth model, led some of its proponents (mainly Pearl) to claim that the model described a *general law* of population growth (Cooper, 1993).

So far, I have made liberal use of the phrase ‘applying to many systems’. But what exactly does ‘applying’ mean? For Levins (1993) and many philosophers of science (e.g. Godfrey-Smith, 2008; Matthewson & Weisberg, 2009; Weisberg, 2013), a model applies to a target when it is used to understand, explain or predict the target system’s dynamics or behaviour. In other words, a model must be able to provide information about the aspects of the system that the scientist considers important for the purpose of the study. This constraint is intentionally weak, so as to allow scientists the freedom to use models in novel and creative ways. For example, it allows scientists to use models developed for quite different systems, such as using physical or chemical models in biology or biological models in economics (Knuuttila & Loettgers, 2016b; MacLeod & Nagatsu, 2016; Railsback & Grimm, 2011; Weisberg,

¹ It is also compatible with, and similar to the standard philosophical notion of generality in the context of laws, where general laws are those that describe general patterns, and are able to subsume many particular phenomena in the world (Kitcher, 1981; see discussion in Mitchell, 2009).

2007). This view is echoed by a number of philosophers of science, who believe that a model's use is partly determined by the intentions of the modeller (Giere, 2004; Godfrey-Smith, 2008; Knuuttila & Loettgers, 2016a; Weisberg, 2013). So, a model *applies* to a target system when a modeller uses it to investigate that target system.

Nonetheless, merely applying a model to a target system does not guarantee that the outcome will be successful. The success of a scientific investigation, whether this is understanding, explaining, predicting or intervening on the target system, also depends on how *well* the model applies to the target. This is where philosophical views begin to diverge, as there are many different accounts of what constitutes *applying well to*, or *fitting*, a target system, such as similarity (Weisberg, 2013), isomorphism (Van Fraassen, 1980), partial isomorphism (Da Costa & Frencha, 2003), mediation (Morgan & Morisson, 1999), fictionalism (Frigg, 2009; Toon, 2012) etc. The aim of this paper is not to provide an argument for any of these accounts; the discussion of applicability and generality presented here is compatible with all views of model-target relations.

There are two aspects of 'applicability' that are relevant for this discussion. The first is being able to identify when a model has been applied well to a system and what factors reduce the applicability of models to targets. A common view in the scientific literature, is that we can *test* that a model applies well to a target, by whether it yields accurate predictions about the system's behaviour (see for example Beckage et al., 2011; Dambacher et al., 2003; Kulmatiski et al., 2011; Levins, 1966, 1993; Stillman et al., 2015; Stockwell, 1999; Tompkins & Veltman, 2006).

The second is that there is a tension between applying *widely* and applying *well*. That is, the more target systems a model is applied to, the less well it tends to apply to each particular target system. This is because in order to generalise, scientists must focus on factors that are common across different systems. In order to achieve this, they omit those factors that represent idiosyncratic features of particular target systems. So, the more general a model, the fewer factors it will have, so as to apply to more systems. The problem is that many systems where general models are used are complex and heterogeneous, so these idiosyncratic factors are actually relevant for the behaviour of each system.² Models without these factors will often yield less accurate results for many particular systems (Berger et al., 2008; Elliott-Graves, 2018; Phillips et al., 2016; Stillman et al., 2015; Travis et al., 2014).

Consider, for example, the *exponential* growth model. It describes the growth of a population in the absence of any factors limiting it, so it is simpler than the logistic growth model.³ It applies to all populations of organisms, because all populations have a growth rate, so it is very general. However, in most cases it does not apply

² A system is *complex* when it has many interacting parts (Levins 1966; Matthewson 2011). A system is *heterogeneous* when these parts are themselves diverse (Elliott-Graves, 2018; Matthewson 2011). A system is causally heterogeneous when the factors that affects its status or behaviour or the dynamics between these factors are idiosyncratic, i.e. particular to that system (Elliott-Graves 2018; forthcoming). Complexity, but especially causal heterogeneity, decrease the transferability of knowledge from one system to another, thus decreasing the scope and extent of generalisations.

³ The exponential growth equation in its differential form is the following: $dN/dt = rN$. Logistic growth is the exponential growth of a population (rN) modulated by the carrying capacity (K).

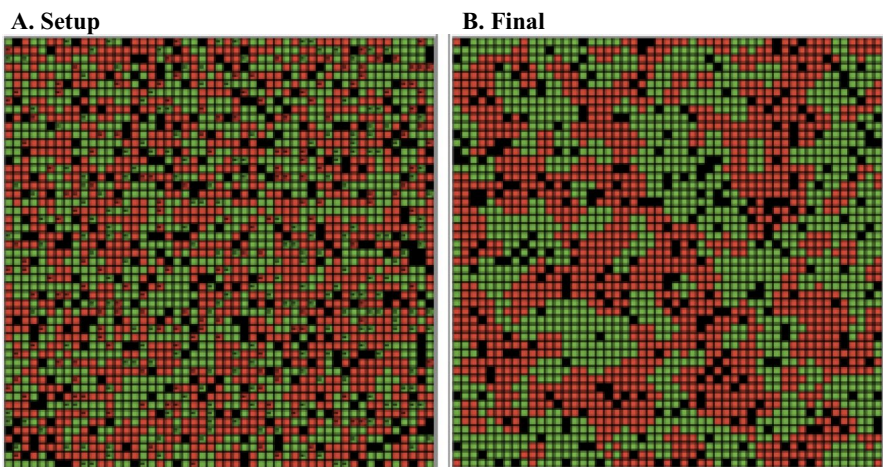
very *well* to biological populations. As stated above, it describes a population's growth in the absence of factors limiting it. Yet most biological populations have many factors limiting their growth, such as scarcity of resources, competition, predation etc. In fact, this model only applies 'well' to populations of bacteria in labs, and only for a limited amount of time, the "logarithmic phase" of bacterial growth (Levins, 1993).

I will return to the implications of this tension between applying widely and applying well in Sect. 5, in the discussion of *overgeneralisation*. In the meantime, I will explain the notion of genericness, exemplified by Schelling's model of residential segregation.

2.2 Segregation

Thomas Schelling's segregation model was intended as an example of how individual behaviour at the micro scale can have large-scale (macro) effects. It is relatively easy to explain segregation by appealing to strong discriminatory preferences (i.e. racism); individuals of race *A* with strong discriminatory preferences will act so as to remove themselves from neighbourhoods populated by those of race *B*. But what if individuals of race *A* express only 'mild discriminatory preferences' (i.e. wanting 30% of their neighbours to be like them)? In Schelling's model, segregation arises even if agents express only mild discriminatory preferences (Schelling, 1969).

In the model, each agent (originally represented by a checker on a checkerboard) can either remain where they are (if their neighbours collectively conform to their preferences) or move to an adjacent square that is closer to their ideal percentage of 'sameness' ratio (neighbours like them). Figure 1 shows a typical NetLogo run of the Schelling model. Figure 1a shows the randomised setup of two populations and



The output of a typical run of Schelling's segregation model in NetLogo using a virtual city that is 51×51 and two types of agents (red and green). All agents prefer to have at least 30% same-colored neighbors.

Fig. 1 Typical NetLogo run of Schelling model for residential segregation

Fig. 1b shows the segregated neighbourhoods after the agents have interacted (segregation usually takes between 15 and 35 ‘ticks’, i.e. agents’ actions).

Unlike the creators of the logistic growth model, Schelling did not base the model on any particular city, nor did he intend it to represent any particular real-world city (Aydinonat, 2007). Rather, it seems that the model was born from Schelling’s need to find an intelligible way to represent an *abstract phenomenon* to his students:

“I decided, if I am going to teach my students, I’ll have to make it all up. I wondered what to do, and decided maybe now was the time to begin playing around with these ideas ... So, I drew a line on a sheet of paper, put down sort of a haphazard ... x’s and 0’s, ... And then I wanted to do it in two dimensions. In one dimension you can simply move an item and insert it between two others. But in two dimensions you have to have a more specific way of deciding where one can go. That is when I decided, well, I could use the checkerboard and leave blank squares so that the movement could get started. (I thought of hexagons, and decided squares were good enough.)”

(Schelling, in Aydinonat, 2005, p. 4)

The model is not supposed to uncover the *actual* mechanism that gives rise to segregation in real-world cities, but an abstract mechanism, which does not provide ‘learning’ about real-world segregation (Fumagalli, 2016). The question is, what is the value of a model that provides knowledge about abstract rather than actual mechanisms of segregation?

Many philosophers of science believe that the Schelling model provides a *how-possibly* explanation of the phenomenon of segregation, by describing a mechanism that *could* give rise to segregation in certain contexts (Aydinonat, 2005; Ylikoski & Aydinonat, 2014). How-possibly explanations are valuable in the following ways: first and foremost, the knowledge that a certain behaviour or effect is possible is valuable irrespective of whether it actually occurs. That is, the model’s results ‘increase the menu of possible explanations’ i.e. bring to light a previously unknown or unexplored mechanism that can cause segregation (Ylikoski & Aydinonat, 2014). This reduces the number of unconceived alternative mechanisms for the segregation, giving us greater confidence in our overall knowledge of the phenomenon.

Second, a how-possibly explanation could provide indirect knowledge for real-world phenomena. Even if the Schelling mechanism has hitherto not been a sufficient cause of segregation, this does not preclude it from constituting a sufficient cause in *another time or place*. Alternatively, the Schelling mechanism could seem to be absent from actual cases of segregation because it was overshadowed by other mechanisms. Yet knowledge of this mechanism could be important for interventions. For example, Schelling mechanisms could result in the persistence of segregation even if the overshadowing factors are eliminated.

Third, even though the Schelling model does not possess *actual* generality (*a*-generality), i.e. it does not apply to many actual systems in the world, it possesses *possible* generality (*p*-generality), i.e. it applies well to many logically *possible* systems (Matthewson & Weisberg, 2009; Weisberg, 2013). The Schelling mechanism that leads to segregation has been shown to be robust across many different modelling scenarios, such as differences in utility functions, neighbourhood sizes,

spatial configurations and rules for updating (Weisberg, 2013 p. 119). *P*-generality is valued for its explanatory power, as “exploration of the non-actual helps explain the actual, and the point of some explanatory models is not necessarily to resemble any real systems, but to canvas possibility space” (Matthewson & Weisberg, 2009, p. 182). In the Schelling case, finding a mechanism that worked in many possible systems could help scientists to determine why it does or does not work in various actual systems.

To sum up, generic models are not meant to be applied to *any* real-world system, let alone *many* real-world systems. Their main value is that they generate how-possibly explanations which can yield information that is useful for real-world systems, but only indirectly.

3 Clarifying the distinction

3.1 Generality, genericness and abstraction

With the examples in place, we can now delve deeper into the difference between generality and genericness. I will start with a short discussion of the term ‘generic’. I have chosen this particular term, because colloquially, ‘generic’ is often used to denote products that do not have a brand name, such as generic ‘cola’. When we say ‘get some cola for the party’, the implication is that we want a certain type of soft drink, but we are indifferent with respect to which brand of cola we want. In other words, the term ‘cola’ refers to an abstract type of soft drink, rather than particular real-world manifestations of it, i.e. actual cans of Coca Cola, Pepsi etc. Similarly, Schelling’s segregation model describes an abstract hypothetical mechanism that could bring about segregation, rather than actual mechanisms that bring about segregation in many real-world cities.

My second reason for picking this term is that it captures one aspect of *abstraction*. The term ‘abstraction’ is notoriously ambiguous in every-day language and in philosophy, yet I believe that disentangling some of this ambiguity can help to elucidate the difference between generality and genericness. In the philosophical literature, abstraction can refer to refer to non-concrete objects, concepts or ideas (Cartwright, 1989), generalizations (Levy, 2018; Strevens, 2008), isolations (Mäki, 2009) or parts of a larger system (Jones, 2005). Whether or not all these terms are coextensive is a matter of ongoing debate (Godfrey-Smith, 2009; Humphreys, 1995; Levy, 2018). A detailed analysis of the entire debate is beyond the scope of this paper, so I will focus on two views that are relevant for the purposes of this discussion, the view of abstraction as non-concreteness and the view of abstraction as generalisation.

Thinking of abstraction in opposition to concreteness has a long history in philosophy and psychology (Rosen, 2017). An important version of this view, in philosophy of science, is Nancy Cartwright’s. Following Aristotle, Cartwright states that in order to abstract, we start off with a concrete particular, an object or system in the world that has all its properties. Then, “we strip away,—in our imagination— all that is irrelevant to the concerns of the moment to focus on some single property or set

of properties ‘as if they were separate’” (Cartwright, 1989, 197). That is, we identify the properties of the object, disentangle them (separate them from each other) and dismiss those that are considered unimportant or irrelevant for the particular model or experiment.⁴ An example of a concrete object is a triangle drawn on a slate. We can abstract by stripping away from the object various properties, such as the colour of the triangle, the chalk that it is drawn from “and other properties incidental to being a triangle” (p. 213). The now abstract system consists of the remaining properties, i.e. the geometrical properties of triangles. As the ‘stripping away’ is done in our imagination, the resulting product is not concrete.

Cartwright’s account of abstraction includes what is usually meant by the term ‘generic’. When we strip actual properties away from a triangle, the resulting imagined triangle ceases to have properties that are particular to actual triangles. The remaining properties *could* but *need not* be applied to any actual triangle. Thus, we could learn about these geometrical properties and their relations, without ever applying them to any concrete triangle. In other words, genericness is the absence of particularity.

I agree with Cartwright that genericness is the absence of particularity, yet I disagree with equating absence of particularity with non-concreteness. Even though stripping away properties in our imagination necessitates leaving the realm of the actual world (to non-concreteness), it does not necessitate genericness. Keeping within the triangle theme, the west pediment of the Parthenon is a concrete triangular object that has a number of other properties: it is made of Pentelic marble, it depicts the contest between Athena and Poseidon for the patronage of Athens, it faces west, and so on. If we focus solely on the triangular aspect of the pediment, we are omitting all the other properties. Nonetheless, if we refer to *that* triangle, we are not referring to a generic triangle, we are referring to a non-concrete but particular triangle without some of its properties. In other words, stripping away is often a prerequisite of genericness but it does not necessarily entail it.

A similar argument can be made with respect to the relationship between abstraction and generality. On some views in philosophy of science, abstraction itself is synonymous with generalisation, because an abstracted system is understood as a general idea that refers to many particularizations of that idea. According to Strevens, for instance, abstraction and generality are one and the same: “of two explanatory models for the same explanandum, one is a better explanation than the other if it is more general, that is, more abstract” (2008, p. 134). On other views, including Cartwright’s, the process of abstraction (i.e. the omission of parts and properties from a system) is a prerequisite of generality. In order to identify what is common across a number of seemingly disparate systems, we must first leave out all

⁴ An important aspect of Cartwright’s view, though not relevant for this discussion, is that this process does not involve distortions. This is what distinguishes abstraction from idealization, where specific properties that are inconvenient are “mentally rearranged”, i.e. distorted (p. 187). Thus, while in the case of idealization a system is misrepresented, a description of an abstract system will contain only truths. It may not contain all the truths, but it will not contain any falsehoods. This point has generated some controversy, as some philosophers argue that abstraction also involves distortions, thus casting doubt on whether abstraction and idealisation are truly separate (e.g. Humphreys 1995; Weisberg 2007).

the factors that are different. What remains is what is common across all systems. Identifying these common factors is the basis for the subsequent generalisation, i.e. the model applies to all the systems that share these common factors.

Yet even if we agree that the process of omission or simplification is a pre-requisite to generalisation, this does not mean that omission entails generalisation.⁵ We can omit or simplify without generalising. Making a statement about the imagined west pediment of the Parthenon, albeit without many of its properties, is not a statement about *all* triangles, or even all the triangles in the Parthenon. It is a statement about a non-concrete single triangle. It is an extra step of extrapolation to make a statement about more (or even all) triangles.

Showing that genericness and generality are both distinct from non-concreteness also helps us avoid another conceptual misstep, namely the idea that *only* generic models are abstract (i.e. non-concrete). The concreteness of the *model* is an orthogonal issue. Non-concrete mathematical models can be general or generic.⁶ The point is what each model is *about*. Whether a model is used generally or generically is determined by what sort of system (or systems) it applies to. Abstraction, qua non-concreteness, is relevant *here*. General models are about many concrete systems, whereas generic models are about generic systems. In other words, it is the *target systems* of the models that differ in terms of concreteness, not the models themselves.

I will provide a more detailed examination of general and generic target systems in Sect. 4. For the remainder of this section, I will address some residual issues that will help to further clarify the difference between generality and genericness.

3.2 Model construction and intended use

As stated in the introduction, my argument places a lot of emphasis on a model's *intended use*. In my discussion of the logistic growth and segregation models, I showed that the former was explicitly intended to apply to many real-world systems, whereas Schelling explicitly did not intend that his model be used to explain segregation in actual cities. The idea that modelers' intentions are important aspects of model construction and use is far from controversial and well established in the literature (Giere, 2004; Knuuttila & Loettgers, 2016a, b; Morgan & Morisson, 1999; Weisberg, 2013). Moreover, modelling is 'outcome-oriented': modelers construct models "keeping an eye on the behaviour they are supposed to exhibit and the results they are expected to produce" (Knuuttila & Loettgers, 2016a, p.1023).

⁵ In fact, omission is not always a prerequisite for generalization. There are instances where scientists start off with an abstract (often mathematical) idea and then apply it to actual systems in the world (see also discussion in Sect. 3.3, and footnote 6). I would like to thank an anonymous reviewer for bringing up this point.

⁶ Indeed, concrete models can also be particular, general or generic. For example, a picture of a plant cell is concrete but generic, when it is not intended to provide information about any actual plant cells. In contrast, a scale model of a particular type of airplane, for use in wind tunnel is general, as it applies to (and provides information about) many actual particular airplanes. Finally, the San Francisco Bay model is concrete and particular, as it is intended to provide information about the actual bay.

However, model use is dynamic, in the sense that a model might be used for different purposes. In some cases, as outlined in Sect. 2.1, models have been applied to systems in entirely different disciplines than the ones for which they were constructed. In other cases, different modelers have arrived at the same model but have had different intentions with respect to its use. For example, Knuuttila and Loettgers (2016a) document the differences between Lotka's and Volterra's construction of the predation model. On their view, Volterra's construction proceeded through isolation and mathematical representation of factors that were present in the real-world system, whereas Lotka's construction started from a generalised mathematical template that was not applied to real world systems (until later).⁷ This latter point is particularly salient, as Lotka's construction of the predation model (which builds on the abstract exponential and logistic growth models) more closely mirrors Schelling's approach in constructing the segregation model than Reed & Pearl's approach to constructing the logistic growth model.

This issue seems to spell trouble for my account. On the one hand, if we are to distinguish between general and generic models, then how do we account for cases where models are used for different purposes, or where generic models become general and vice versa? If, on the other hand, we agree that models can be used for different purposes, then what is the point of making the distinction in the first place?

The short answer is that models are neither general or generic in isolation – they are used as general or as generic models. If models were always either general or generic, the distinction would be obvious, and could be easily demonstrated just by the two examples in Sect. 2. The rest of this discussion would be unnecessary. The point is that the possibility of changing a model's use does not mean that anything goes. In other words, just because it is possible to change a model's use, this does not mean that all such changes are equally *useful*, worth implementing or will lead to successful results. Acknowledging and highlighting the distinction can help us understand when we can switch between genericness and generality and when we should not.

First, the switch from generality to genericness is less fraught with dangers than the opposite. This is because if we have already established that a model is general, we have identified a mechanism, dynamic, or set of causal factors that are shared across many real-world systems. Once such a mechanism is known, then it is often worth to explore the space of theoretical possibilities so as to gain knowledge about hypothetical systems that are adjacent (in possibility space) to the actual systems we have already investigated. For example, a general model of sexual reproduction can be turned into a generic model of reproduction to examine how organisms with 3 or 4 sexes *could* reproduce (see Sects. 4, 5.1 and Weisberg, 2013, for a discussion of how to interpret Crow's model of sexual reproduction in this context). In Sect. 5.1,

⁷ I should note that this point reveals a limitation of Cartwright's account of abstraction, in the sense that it does not accurately capture the actual way in which all models are generated. While scientists sometimes start off with phenomena in the world and strip away properties to arrive at abstract models, in other cases scientists start off with abstract ideas or concepts and construct models without any reference to real-world phenomena. I would like to thank an anonymous reviewer for pointing this out.

I will show that the main danger for general models is not that they are (relatively) easily transformed into generic models, but that lumping together generality and genericness creates expectations of generality that are beyond the scope of many general models.

Second, there are differences in the type of system that models apply to, which facilitate the (relatively rare) cases where generic models can be used as general models. Some *systems* lend themselves to both generic and general model use. Often, these systems are physical or chemical systems, which are complex but homogeneous, i.e. the relevant causal factors are shared across systems, while any idiosyncratic features of a particular system or any differences between them are irrelevant and can therefore be safely ignored (Elliott-Graves, 2018; Knuuttila & Loettgers, 2016b, Weisberg, 2013, see also footnote 1). A good example of this is the Ising model, which can be used both generally and generically. The Ising model describes phase transitions, where a system transitions from one state into another. It was originally developed to study phase transitions in ferromagnetic materials: when the temperature of a ferromagnetic system rises above a ‘critical’ point T_C , the system becomes paramagnetic (Knuuttila & Loettgers, 2016b). Knuuttila & Loettgers argue that the Ising model should be viewed as a *model template*, an abstract conceptual idea, expressed mathematically or computationally.⁸ In this sense, the Ising model is being used generically, as it describes the abstract phenomenon of phase transition, without reference to particular (micro-level) properties of any actual system. However, it has since been shown to apply to many actual systems in physics and chemistry, in each case describing the phase transition of the system in question (ibid). It can also describe situations as mundane as boiling water, and its subsequent change from a liquid to a gaseous state.

What are we to make of cases like the Ising model? I believe that the correct response is to be very happy that it was possible for a generic model to become general. I believe scientists should rejoice in the insights that the Ising model generated. Nonetheless, we should also exercise caution, for two reasons. First, just because such examples are possible, does not mean that we should expect generic models to become generic *often*. It just so happens that in these physical and chemical systems, there is sufficient homogeneity within and across systems so that the generic idea manifests in sufficiently similar ways in many actual systems so that the model can apply to them. We should not expect such transformations to be frequent or particularly useful, especially when the systems under investigation are complex and heterogeneous.⁹

This leads to the second reason to exercise caution. Even though a generic model can apply to a homogeneous group of systems, this does not mean that it will

⁸ In fact, I believe that my distinction can allay the worry expressed by Batterman and Rice (2014), who point out that “common features accounts” do not fully capture the importance or source of value of models like the Ising model, as some of its value is due to its ability to explain the dynamics of a highly abstracted phenomenon. Though I believe that they underplay the importance of generality, I agree with them that genericness is valuable independent of its relation to generality.

⁹ See footnote 2 and Sect. 5.1 for a discussion of heterogeneity.

apply to other types of systems. Socio-physicists (a small group of physicists who attempted to give physical explanations of socio-economic phenomena) recognised a “physical analogue” between the Ising model and the ‘clustering’ observed in Schelling’s segregation model (Vinković & Kirman, 2006 in Knuuttila & Loettgers, 2016b, p. 391). They argued that it is possible to conceptualise the abstract phenomenon of segregation a phase change where the “system is frozen into a solid” (ibid). In generic terms, this is both interesting and insightful. As Knuuttila & Loettgers point out, the Ising model shows “what happens if neighbourhood segregation or opinion formation are understood as some kind of herding” (p. 397). Such a result is valuable because “much of economic modelling can be understood in terms of conceptual exploration”. However, assuming that this means the model will apply to actual cities is premature, as we should not expect that the Ising model applied to social systems will have “the same predictive value, and empirical and theoretical grounding as in physics” (ibid). I agree with this analysis, as I believe that it shows that the Ising model can be used generally within physics (and chemistry) yet is unlikely to yield useful information when applied to particular socio-economic systems. In these cases, it is better to use the model generically (see Sect. 5).

To sum up the argument in this section, it is precisely because it is sometimes possible to change a model’s intended use that we should be careful, even cautious, in implementing such changes. I will return to the topic of intended use in Sect. 4, where I show that it is also important in determining the target of a model and in Sect. 5, where I show there are serious epistemic reasons to resist the use of the generic Schelling model as a general model, at least without *additional* empirical tests.

3.3 Model families, types, templates & uninstantiated models

Most models are not completely unique, but can be grouped together with other related models, in so-called ‘model families’. For example, the family of growth models (which include the logistic and exponential growth models) can be written as $dN/dt = \alpha N + \beta N^{1+\alpha}$ (Evans et al., 2013; Weisberg, 2013). Another way to organise or classify similar models is in terms of model types or templates. Model types can be thought of as sets, in which individual models are members, while a model template is “a mathematical structure that is coupled with a general conceptual idea that is capable of taking on various kinds of interpretations”, that is, an abstract conceptual idea that is associated with certain mathematical forms and computational methods (Knuuttila and Loettgers 2016b p. 396).

Each of these terms correlates with certain intuitions. For instance, we might think of families of models as general, in the sense that they encompass many particular models. On the other hand, we might think of model templates as generic, because a template is an abstract (in the sense of non-specific) version of a number of models. To make matters worse, as outlined in the previous section, cases such as the Ising model can be thought of as general or generic. Do these cases undermine my distinction?

Again, the answer is no. I have argued that models themselves are neither general nor generic, but can be used in general or generic ways. The same point applies

to model families, types, or templates. These model groupings can be used in both ways: generally, if they are used to provide information for many particular systems and generically, if they are used to provide information about abstract or hypothetical systems. Still, I expect these model groupings to be used generically more often than they are used generally. We can think of model groupings as mapping a set of trajectory spaces, whereas each particular model describes the set of trajectories within a particular space. The information provided by the model groupings is useful for understanding the scope and features of the state space, whereas it is the individual models that provide information about the actual trajectories of real-world systems. In other words, the information yielded by the model grouping is more likely to apply to abstract, hypothetical versions of real-world systems.

Similar points can also be made with respect to uninstantiated models, i.e. those that do not have specified values for their parameters. In their critique of Levins, Orzack and Sober (1993) argue that uninstantiated models are, by definition, more general than instantiated models. This is a rather odd claim for them to make, since they subsequently point out that we should not compare instantiated and uninstantiated models: “Comparing apples with apples and oranges with oranges requires that the two both be treated consistently as uninstantiated or instantiated” (p. 536). Moreover, Orzack and Sober accept that generality is defined as applying to many actual systems, yet uninstantiated models can only apply to actual systems if they are first instantiated. Still, a more charitable interpretation of the argument is the following: a model with specified parameters describes a complete set of trajectories through a state space, where the dimensions are the variables of the equation. An uninstantiated model corresponds to a set of trajectory spaces. Nonetheless, the problem is that this set of trajectory spaces does not pick out more systems *in the world* – it picks out more hypothetical systems. Thus, it provides direct information about the space of possibilities and only provides indirect information about actual systems. In short, a model that describes a set of trajectory spaces is generic rather than general.

3.4 Source of value

Another difference between general and generic model uses concerns the source of their value. Generality is valuable because it allows scientists to formulate expectations about a phenomenon or system even if they do not directly investigate or manipulate it. Thus, identifying that a particular phenomenon is a manifestation of a general pattern means that scientists can compare its behaviour to the other similar systems. It also allows scientists to make predictions about how the system will behave in the future, based on the behaviour of these other similar systems.

When models are used generically, they do not apply to any systems in the world, so generality cannot be the source of their value. Instead, it is their ‘genericness’ that makes them valuable as it provides scientists with hypothetical situations that can be used to explore possibilities, unencumbered by constraints imposed by the real world. Thus, scientists can determine how certain situations *could* come about, by exploring what conditions need to be in place for them to occur, and which

conditions constitute limiting factors. Of course, this process also yields information about the actual world, indirectly, because it helps scientists gain understanding about how these factors and conditions function and interact with each other. Still, this is a secondary source of value. Using a model generically is valuable primarily because it allows scientists to explore the space of possibility, and this value remains unaffected by any subsequent discovery that they can be applied to real-world phenomena, or indeed by how many real-world phenomena they can be applied to. It is useful to know how a perpetual motion machine would work if it existed, or that entities can make flock-like shapes when flying, based on a single behavioural rule, or that a micro-behaviour can have a macro effect. *How* useful this knowledge is, compared to the knowledge gained by investigating other types of models, is beside the point (but see Weisberg, 2013 for a discussion of this). The point is that it *is* useful, in itself.

3.5 A note on a- and p-generality

It may seem that the distinction between generality and genericness nothing more than the distinction between actual (*a*-) versus possible (*p*-) generality (Matthewson & Weisberg, 2009). After all, models used generally have high levels of *a*-generality (they apply to many actual systems) and models used generically have high levels of *p*-generality (they apply to many possible systems). However, there are three important differences between the two distinctions. The first is that the *a*- vs. *p*- distinction does not map on to the two kinds of models cleanly, as many models have high *p*-generality *in addition* to having high *a*-generality. For instance, the logistic growth model applies to many actual and non-actual systems, including computer populations, populations that have yet to be developed in labs, populations of species that are extinct and populations of species that have yet to evolve.

Second, in the case of generic model use, having high *p*-generality is *independent* of and *orthogonal* to being generic. Possessing *p*-generality does not preclude models used generically from having other characteristics. Models can be applied to a generic phenomenon *and* many particular *hypothetical* phenomena. For example, the Schelling model explains the generic phenomenon of segregation, but it can also be applied to (and predict the outcome of) segregation in a hypothetical ‘neighbourhood’ in NetLogo. The subsequent application of the model to the hypothetical system does not negate or change the fact that the model was originally intended to represent an abstract phenomenon.

Third, possessing *p*-generality is also not the only reason to value models used generically. In the previous section I argued that the value of the Schelling mechanism is not affected by how many actual cities it applies to. Yet this also applies to *p*-generality. The value of the Schelling mechanism is not entirely determined by how robust it is across various hypothetical systems. Even if it only came about in some circumstances, or within certain parameters, that would still constitute an important insight for how micro-level behaviours can have macro-level effects. In fact, high levels of *p*-generality can sometimes detract from the value of the model. In generic models, high *p*-generality can indicate that the model’s output

is an artefact of the model's design. For example, Muldoon et al., (2012) ran versions of Schelling's model and found that segregation emerged even in cases where individual agents preferentially chose neighbours unlike themselves. Their analysis revealed that "as long as agents care about the characteristics of their wider community, they tend to end up in a segregated state" (p. 38), which implies that there is something about the setup of the model and the how agents process information that causes segregation in this case. The worry then is that these (or other) effects of the particular game might be affecting or even driving a particular output. Thus, focusing exclusively on p-generality fails to recognize an important source of value for generic model use.

So far, I have argued that there are important differences in the intended use and source of value between general and generic model use. In the next section, I argue for a third difference between them, namely that general and generic model use involve different types of target systems.

4 Targets of general models

If we are to understand what general models are *about*, then we need to examine their target systems. 'Target systems' are important components of model-world relations, and there is a consensus in the literature that they are what models represent. Yet 'target systems' can be understood in a number of different ways, e.g. as parts of real-world systems, abstract systems, or even hypothetical systems, depending on the types of models they are associated with and the use for which they are intended (Elliott-Graves, 2020; Giere, 2004; Knuuttila, 2005; Weisberg, 2013). Instead of going through all these accounts of targets, I will focus on Weisberg's (2013) account, as it is the most comprehensive discussion of general models and their targets in the literature.

For Weisberg, general models do not target any particular manifestation of a phenomenon in a real world-population, but apply to a *generalised* target, which is an abstraction over many particular targets (Weisberg, 2013, p. 116). Weisberg illustrates his view with an example from biology, the evolution of sexual reproduction. He points out that the emergence of sex is a strange phenomenon and there is no real consensus on why it evolved. Our best bet for determining why sexual reproduction confers some evolutionary advantage is constructing and analysing a generalized model of sexual reproduction. In other words, the point of studying sex in a general manner is not to discover facts about sexual reproduction in particular populations, but to understand larger issues such as the relative merits of sexual reproduction when compared to asexual reproduction. A "generalized model of sexual reproduction isn't supposed to be about kangaroo sex or fungi sex, but about sex itself." (p. 115). Consequently, on Weisberg's view, the target of a model of 'sex itself' will also be generalised.

Yet Weisberg points out that there is no such thing as 'sex itself' in the world, only a collection of instances of sex in particular populations "kangaroo sex, Tasmanian devil sex, human sex, but not sex in general" (p. 116). Still, all these sexually reproducing species have certain properties in common. These shared properties are the

core causal features of the general phenomenon of sexual reproduction. So, if we find the *intersection between the various instances of sex*, we can create a system which the general model will represent. In other words, “sex in general is an abstraction over these more specific kinds of sex” (p.116). The general model targets this generalised system which is an abstraction over the many actual targets of sexual reproduction.

The problem with this picture is that Weisberg lumps together aspects of modeling that ought to be kept distinct, which leads to a misrepresentation of both general and generic model use. Weisberg states that when scientists use general models, they are not interested in learning about particular manifestations a phenomenon, such as sexual reproduction, but about the phenomenon in general. Yet this seems to be at odds with actual scientific practice. There are numerous cases where scientists use models generally, i.e. with the explicit intent to learn about many real-world manifestations of that phenomenon. I documented at least one instance of such intentions in Sect. 2.1, in my discussion of the logistic growth model. There are many other examples of models being used to learn about many different systems in the world. Just within biology, for example, niche models (Sobek-Swant et al., 2012; Wang & Jackson, 2014), loop analysis models (Dambacher et al., 2003; Puccia & Levins, 1986), individual-based models for a number of different dynamics (Berger et al., 2008; Railsback & Grimm, 2011) drought sensitivity models (Phillips et al., 2016), and fisheries models (Hollowed et al., 2000) are just some examples of models explicitly intended to apply to many real-world systems.

The notion of a ‘generalised target’ does not make sense in the context of general models. Recall that this is an abstract generalisation of the intersection of many real-world targets. In Weisberg’s example, “the theorist takes the features common to all instances of sexual reproduction to be the target of her model.” (p. 116). But if a scientist can create an intersection from many real-world targets, then the generalised targets are dependent on specific targets. In other words, using the term ‘generalized target’ is simply shorthand for referring to a family of actual targets which share common properties, so the model is about kangaroo sex *and* fungi sex *and* Tasmanian devil sex and so on. If the intended use of the model is to apply to many real-world systems, then the whole process of finding the intersection is superfluous, so there is no need to invoke the notion of a ‘generalised target’.

Nor does the ‘generalised target’ accurately describe how targets of models intended to be generic are generated. As we saw in Sect. 2.2, Schelling did not in fact assemble various instances of segregation in real cities, identify the intersection between them and create an abstract system of that intersection. He approached the issue from a different perspective, coming up with a mechanism that led to segregation in an abstract, hypothetical system. He did not even test whether the model explained segregation in actual cities. So, it seems rather odd to attribute this bottom-up strategy intersection strategy to Schelling, or indeed any other instance of generic modelling. If scientists do not wish to learn about how a phenomenon manifests in the world, then why create a target that is constrained by how the phenomenon does manifest in the world?

Would it not be easier to approach the topic in a hypothetical, non-specific context? Indeed, Weisberg outlines exactly such a context in his discussion of hypothetical models, which are a much better description of Schelling’s approach (Weisberg, 2013,

pp. 121–122). Targets of hypothetical models, such as frictionless planes or perpetual motion machines do not or cannot exist. The similarities with models used generically are striking, as their value is to generate how-possibly explanations, these models and their hypothetical targets can inform us about the world, though they do so indirectly. For example, learning about the behaviour of a frictionless oscillator can help us to understand the movement of similar systems with friction, such as pendula.

To sum up, models used generally have many local targets. Here, the notion of a generalized target is unnecessary, or at least merely shorthand, as it is just a way of referring to all the factors that are common across the many local targets. In contrast, in the context of generic model use, scientists construct a model that does not refer to any actual targets, but a *generic* target. This target will have the same ontological status as a hypothetical target and will have the features that bring about the phenomenon being investigated *in a hypothetical system*.

5 (Mis)application of general models

In the previous sections, I distinguished between general and generic model use on largely theoretical grounds, i.e., by arguing that there is a conceptual difference between generality and genericness with respect to what models are about. Yet the importance of this distinction is not merely theoretical, as it has implications for how models are used in scientific practice. More specifically, failing to distinguish between the two conceptions of generality affects their success. That is, there are contexts in which models can be *misapplied*, leading to overgeneralisation or inappropriate application to actual systems.

5.1 Overgeneralisation

In Sect. 3.1., I argued that when used generally, models face a tradeoff between applying *widely* and applying *well*. Models attain generality by omitting factors that are idiosyncratic to particular systems, including only those factors that are common across different systems. However, in complex and heterogeneous systems the idiosyncrasies are real causal factors, not mere noise, so omitting them can lead to predictive inaccuracy. The problem is that in many disciplines there is motivation, even pressure, to produce generalisations (Elliott-Graves, 2019).

This pressure is supported by traditional views in philosophy of science, where generalisations formed the basis for scientific explanation and prediction, which in turn, *constituted* scientific theorising. On many influential accounts, to provide a scientific explanation was synonymous with demonstrating how a particular phenomenon is an instance of a more general pattern, while scientific theories were those that subsumed many disparate phenomena under one framework (Cartwright, 1989; Hempel & Oppenheim, 1948; Kitcher, 1981). The simpler the theoretical framework (in terms of the number of explanatory principles and ‘brute facts’ it appealed to) and the greater the number of phenomena it could explain, the more explanatorily powerful it was.

How extensive is overgeneralisation and what is its importance? The answer depends on the type of investigation or discipline, but there are some disciplines where overgeneralisation is quite extensive. A pertinent set of examples comes from history of ecology. The quest for generalisation in ecology is arguably as old as the discipline itself. For some, ecology only became a truly scientific discipline when it applied simple models from physics to ecological populations (Kingsland, 1995). Before these models were incorporated into the field, ecology was simply a collection of observations of nature (Shrader-Frechette & McCoy, 1993). The quest for generalisation continues as researchers search for general theories that will unify disparate ecological phenomena (Borer et al., 2014; Scheiner & Willig, 2008). However, ecological research is also riddled with exceptions, surprises and unpredictability (Beckage et al., 2011; Kaunisto et al., 2016; Peters, 1991). The cause of this, to a large extent, is the fact that models (used generally) do not take into account the idiosyncrasies of particular systems, so their results are often inaccurate (see Elliott-Graves, *forthcoming* for an extended discussion of this point and Phillips et al., 2016, for a particular example).

Some ecologists are highlighting the idiosyncratic nature of ecological systems and advocating a more modest approach to generalisations (see for example, Beckage et al., 2011; Kaunisto et al., 2016; Phillips et al., 2016). However, these attempts are often criticised (sometimes in high-ranking journals), for focusing too much on data-rich models and failing to see the bigger picture, i.e. how seemingly disparate phenomena can be subsumed under general patterns or laws (Houlahan et al., 2017; Ward et al., 2014).

The problem of overgeneralization applies to models when they are used generally, because this is the context in which there is a tradeoff between applying widely and applying well. Generic model use does not incur a tradeoff. However, failing to distinguish between generality and genericness can lead researchers to underestimate the importance of the tradeoff and overestimate the scope of generalizability in the context of general model use. This is exacerbated if we think that models used generally have generic targets. Recall Weisberg's account of the model of sexual reproduction, which "isn't supposed to be about kangaroo sex or fungi sex, but about sex itself. *"This is why Crow's model of the origin of sexual reproduction has no detail about any specific species in it"* (p. 116, my emphasis). The assumption is that the particularities of each system are mere details that need not be included in the model's target. When all the idiosyncrasies are left out, each local target is identical to the generic target. This can lead researchers to expect that when a model applies to the generic target, it will apply equally well to many local targets. Yet as we have seen, this assumption cannot be made when a model is applied to many complex and heterogenous systems.

5.2 The danger of moving from genericness to generality


This leads to the second issue, namely the dangers of moving from genericness to generality. That is, when models that are not meant to be about any systems in the world are applied to particular systems in the world, because of their perceived

generality. As we saw in Sect. 3.3, this movement is possible in specific cases (such as the Ising model), yet even these cases are limited. The Schelling model is an example of a model that should not be used generally, though as we shall see, it is quite easy to succumb to this temptation.

Weisberg's analysis shows exactly how easy it is. Weisberg claims that "one way to use Schelling's model is to try to account for the patterns of segregation in a real city" (p. 118). He starts by observing that the Schelling model's results have been shown to be robust over a number of different modelling scenarios, including different "utility functions, different rules for updating, differing neighbourhood sizes, and different spatial configurations" and that "[i]n fact, it is extremely hard to avoid segregation when agents have some preference for like neighbors" (Weisberg, 2013, p. 118). Then, he states:

For example, in Figure 7.1 [reprinted here as Fig. 2], I show the population density of African Americans in Philadelphia's census tracts next to an example output of a Schelling model. Both have a 75% exposure index, meaning that 75% of neighbors are of the same race. *This use of Schelling's model is an example of target-directed modeling.*" (p. 118, my emphasis)

While Weisberg does not explicitly state that this is an actual explanation of segregation in Philadelphia, his analysis implies that the application of the Schelling model to actual systems is a legitimate instance of target-directed modelling. But is it legitimate? What we know is that the Schelling mechanism *could* cause segregation, and that the segregation outcome is robust *in various modelling scenarios*. We also know that Philadelphia neighbourhoods have a similar segregation ratio. Is this enough to infer that the *actual* cause of segregation in Philadelphia was the Schelling mechanism? I do not believe that it is. The Schelling model provides information



The output of a typical run of Schelling's segregation model using a virtual city that is 51×51 and two types of agents (gray and white). All agents prefer to have at least 30% same-colored neighbors.

Patterns of racial segregation in Philadelphia from the 2010 census. Darker areas correspond to higher percentages of African-Americans

Fig. 2 Schelling's model applied to Philadelphia (reprinted from Weisberg, 2013 p. 119)

about a mechanism of segregation, yet in order to determine whether this is the actual mechanism that explains segregation in Philadelphia, we must first consider and rule out alternative causes or mechanisms of segregation.¹⁰

These alternative causes/mechanisms exist and are well known! In Sociology, segregation is explained through alternative mechanisms, i.e., deliberate and sustained governmental policies such as redlining (Rothstein, 2017; Sampson, 2019). For example, some scholars believe that individual preferences *could not* have caused the type of segregation that exists in Philadelphia: “Today’s residential segregation ... is *not the unintended consequence of individual choices* and of otherwise well-meaning law or regulation but of unhidden public policy that explicitly segregated every metropolitan area in the United States. The policy was so systematic and forceful that its effects endure to the present time.” (Rothstein, 2017 p. viii, my emphasis).

What should we make of this? At the very least, I believe that it shows that the Schelling mechanism cannot be assumed to explain segregation in real cities, without further investigation. That is, before we can claim that the Schelling mechanism is what caused segregation in Philadelphia, we must determine that the alternative factors *did not cause it*. It may turn out that the Schelling mechanism did cause segregation, that it did not, or that it contributed to segregation in actual cities. The point is that these outcomes are yet to be determined. In fact, Fumagalli (2016), argued that the output of the Schelling model should not be taken to undermine or reduce confidence in the belief that segregation is caused by explicitly racial preferences. It could provide learning about real-world instances of segregation, if supplemented with “information or presuppositions regarding those real-world situations” (p. 449), but then, as Fumagalli points out, it would cease to be an abstract model.

The point can be illustrated further by contrasting the generic target system with a group of local target systems. As shown in Sect. 2.2., Schelling did not intend his model to apply to any real-world systems, but to a non-specific, aka (generic) system. It could be the case that the generic target bears some similarity to many local targets. That is, it could be the case that mild discriminatory preferences are present in real-world cities and that these preferences induce people to move in and out of neighbourhoods. It could even be the case that these mild discriminatory preferences are the only common factor across all the local targets. Yet even if this is the case, this causal factor might not be sufficient to cause segregation. We must determine that the other causal factors are not the actual cause of segregation, even if they are idiosyncratic (i.e. only present in one or a few real-world systems).

Consider again the case of exponential growth. It applies to all populations, as they all would grow exponentially in the absence of factors limiting them. But all populations do have other factors limiting them, hence actual population growth is

¹⁰ Another interpretation of Weisberg’s claim could be that the Schelling mechanism applies to Philadelphia, though it only provides a how-possibly explanation of the mechanism of segregation in that city. While it is true that the Schelling mechanism provides a how-possibly explanation of segregation in Philadelphia, we do not need to apply the model to the actual city to gain this information. The model provides a how-possibly explanation because it is generic, i.e. it explains how segregation comes about in hypothetical cities. The problem is, however, that it is more likely than not that alternative mechanisms are the actual cause of segregation in Philadelphia. Thus, applying the generic model to the actual city will either give us information that we already have or information that is probably false.

explained by the exponential dynamic and the various limiting factors. These factors can be different in each case (e.g. density dependence, competition, predation, pathogens etc.). While the exponential dynamic is part of the explanation or prediction, the inclusion of other factors is essential for the model's results to be accurate. Yet the limitations of the exponential growth model are well known. Scientists do not use the exponential growth model to explain or predict the growth of actual populations.

The same caution should be employed when applying the Schelling model to actual cities. Even if it turns out that mild discriminatory preferences exist in actual populations, they might be overshadowed by a number of other causal factors. Those causal factors (possibly in addition to the Schelling mechanism) will explain the phenomenon of segregation in actual cities. In other words, in order for the generic target to be sufficiently similar to many local targets (so that the model can be used generally rather than generically) the local targets must include no other causal factors that overshadow the Schelling mechanism. While this is theoretically possible, I believe it is unlikely.

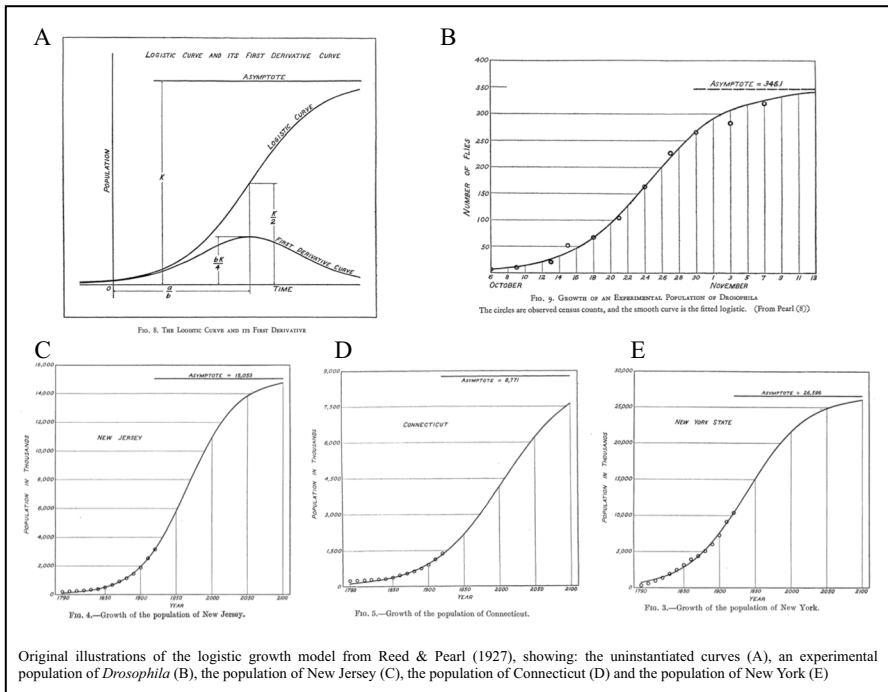
In fact, the Schelling case is an especially good example of how dangerous it can be to use models meant for generic targets as though they apply to many actual target systems. One of the implications of the Schelling mechanism is that segregation emerges even when people are not explicitly racist, as 'mild discriminatory preferences' are contrasted to 'strong discriminatory preferences' and only the latter constitute true racism (see for example Aydinonat, 2007; Weisberg, 2013; Ylikoski & Aydinonat, 2014).¹¹ It would be nice if segregation was not anyone's fault, but just happened to emerge in actual cities from mechanisms that are mild (aka understandable, or at least, not overly blameworthy). However, it would also be wrong to explain actual instances of segregation in this way if it is in fact caused by explicit racism, at the individual or institutional level. I am not accusing any of the authors mentioned in this paper of doing anything wrong. However, I believe that more attention should be paid to the alternative causes of segregation identified by sociologists and economists. It would be very interesting to actually test whether the Schelling mechanism has caused or contributed to segregation in actual cities. Until these tests are performed, however, philosophers should employ even more care not to imply that the Schelling model applies to actual cities.

6 Conclusion

Scientists can use models for a number of different purposes, such as explaining and predicting particular phenomena in the world or acquiring knowledge about generic phenomena that do not manifest in the world in the same way. Both these purposes are legitimate and important, yet they should be kept distinct. Moreover, while models can sometimes be used for both purposes, we should be cognisant of the limitations and difficulties associate with each purpose. Thus, when used generally, models face a tradeoff between applying widely and applying well, so their generality is limited. This is especially important when applying a model intended for generic use to explain or predict one or more phenomena in the real-world.

¹¹ My view is that the term 'racism' is applicable to both types of discriminatory preference, i.e. stronger or milder racism, but that is an issue for another paper.

Appendix



Box 1 Original illustrations of logistic growth – reprinted from (Reed & Pearl, 1927)

Acknowledgements I would like to thank Luis Mireles-Flores and Michael Weisberg, along with the anonymous reviewers and the audiences at the University of Helsinki and the Hebrew University of Jerusalem for their helpful and constructive comments.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest No.

Ethical approval Not applicable.

Informed consent Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aydinonat, N. E. (2005). An interview with Thomas C. Schelling: Interpretation of game theory and the checkerboard model. *Economics Bulletin*, 2(2), 1–7.
- Aydinonat, N. E. (2007). Models, conjectures and exploration: An analysis of Schelling's checkerboard model of residential segregation. *Journal of Economic Methodology*, 14(4), 429–454.
- Barlow, J. (1992). Nonlinear and Logistic Growth In Experimental Populations of Guppies. *Ecology*, 73(3), 941. <https://doi.org/10.2307/1940170>
- Batterman, R. W., & Rice, C. C. (2014). Minimal model explanations. *Philosophy of Science*, 81(3), 349–376.
- Beckage, B., Gross, L. J., & Kauffman, S. (2011). The limits to prediction in ecological systems. *Ecosphere*, 2(11), art125. <https://doi.org/10.1890/ES11-00211.1>
- Berger, U., Piou, C., Schiffers, K., & Grimm, V. (2008). Competition among plants: Concepts, individual-based modelling approaches, and a proposal for a future research strategy. *Perspectives in Plant Ecology, Evolution and Systematics*, 9(3–4), 121–135. <https://doi.org/10.1016/j.ppees.2007.11.002>
- Brashares, J. S., Werner, J. R., & Sinclair, A. R. E. (2010). 'Social meltdown' in the demise of an island endemic: Allee effects and the Vancouver Island marmot. *Journal of Animal Ecology*, 79(5), 967–953. <https://doi.org/10.1111/j.1365-2656.2010.01711.x>
- Borer, E., Harpole, S., Adler, P., Lind, E., Orrock, J., Seabloom, E., & Smith, M. (2014). Finding Generality in Ecology: A Model for Globally Distributed Experiments. *Methods in Ecology and Evolution*, 5(1), 65–73.
- Cartwright, N. (1989). *Nature's Capacities and Their Measurement*. Oxford University Press.
- Cooper, G. (1993). The competition controversy in community ecology. *Biology and Philosophy*, 8(4), 359–384.
- Cooper, G. (1998). Generalizations in Ecology: A Philosophical Taxonomy. *Biology & Philosophy*, 13(4), 555–586. <https://doi.org/10.1023/A:1006508101996>
- Da Costa, N. C., & French, S. (2003). *Science and partial truth: A unitary approach to models and scientific reasoning*. Oxford University Press on Demand.
- Dambacher, J. M., Li, H. W., & Rossignol, P. A. (2003). Qualitative predictions in model ecosystems. *Ecological Modelling*, 161(1–2), 79–93. [https://doi.org/10.1016/S0304-3800\(02\)00295-8](https://doi.org/10.1016/S0304-3800(02)00295-8)
- Douglas, H. (2009). Reintroducing prediction to explanation. *Philosophy of Science*, 76(4), 444–463. <https://doi.org/10.1086/648111>
- Elliott-Graves, A. (2018). Generality and causal interdependence in ecology. *Philosophy of Science*, 85(5), 1102–1114.
- Elliott-Graves, A. (2019). The future of predictive ecology. *Philosophical Topics*, 47(1), 65–82
- Elliott-Graves, A. (2020). What is a target system?. *Biology & Philosophy* 35(2), 28. <https://doi.org/10.1007/s10539-020-09745-3>
- Elliott-Graves, A. (forthcoming). *Ecological complexity*. Cambridge University Press.
- Evans, M. R., Grimm, V., Johst, K., Knuuttila, T., de Langhe, R., Lessells, C. M., Merz, M., O'Malley, M. A., Orzack, S. H., & Weisberg, M. (2013). Do simple models lead to generality in ecology? *Trends in Ecology & Evolution*, 28(10), 578–583.
- Frigg, R. (2009). Models and fiction. *Synthese*, 172(2), 251–268. <https://doi.org/10.1007/s11229-009-9505-0>
- Fumagalli, R. (2016). Why we cannot learn from minimal models. *Erkenntnis*, 81(3), 433–455.
- Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, 71(5), 742–752.
- Godfrey-Smith, P. (2008). Models and fictions in science. *Philosophical Studies*, 143(1), 101–116. <https://doi.org/10.1007/s11098-008-9313-2>
- Hempel, C. G., & Oppenheim, P. (1948). Studies in the Logic of Explanation. *Philosophy of Science*, 15(2), 135–175. <https://doi.org/10.2307/185169?ref=search-gateway:0c1dbed607f660b0c1569ddf721ed555>
- Hollowed, A., Bax, N., Beamish, R., Collie, J., Fogarty, M., Livingston, P., Pope, J., & Rice, J. C. (2000). Are multispecies models an improvement on single-species models for measuring fishing impacts on marine ecosystems? *ICES Journal of Marine Science*, 57(3), 707–719. <https://doi.org/10.1006/jmsc.2000.0734>
- Houlahan, J., McKinney, S., Anderson, M., & McGill, B. (2017). The priority of prediction in ecological understanding. *Oikos*, 126(1), 1–7.
- Humphreys, P. (1995). Abstract and concrete. *Philosophy and Phenomenological Research*, 55(1), 157–161.

- Jones, M. R. (2005). *Idealization and abstraction: A framework*. In *Idealization XII: Correcting the model* (pp. 173–217). Brill.
- Kaunisto, S., Ferguson, L. V., & Sinclair, B. J. (2016). Can we predict the effects of multiple stressors on insects in a changing climate? *Current Opinion in Insect Science*, 17, 55–61.
- Kingsland, S. (1982). The refractory model: The logistic curve and the history of population ecology. *The Quarterly Review of Biology*, 57(1), 29–52.
- Kingsland, S. (1995). *Modeling Nature*. University of Chicago Press.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science*, 48(4), 507–531. <https://doi.org/10.2307/186834>
- Knuuttila, T. (2005). Models, representation, and mediation. *Philosophy of Science*, 72(5), 1260–1271.
- Knuuttila, T., & Loettgers, A. (2016a). Modelling as Indirect Representation? The Lotka–Volterra Model Revisited. *British Journal for the Philosophy of Science*, axv055–30. <https://doi.org/10.1093/bjps/axv055>
- Knuuttila, T., & Loettgers, A. (2016b). Model templates within and between disciplines: From magnets to gases – and socio-economic systems. *European Journal for Philosophy of Science*, 6(3), 377–400. <https://doi.org/10.1007/s13194-016-0145-1>
- Kulmatiski, A., Heavilin, J., & Beard, K. H. (2011). Testing predictions of a three-species plant–soil feedback model. *Journal of Ecology*, 99(2), 542–550. <https://doi.org/10.1111/j.1365-2745.2010.01784.x>
- Kuorikoski, J., & Lehtinen, A. (2009). Incredible Worlds, Credible Results. *Erkenntnis*, 70(1), 119–131. <https://doi.org/10.1007/s10670-008-9140-z>
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54(4), 421–431. <https://doi.org/10.2307/27836590>
- Levins, R. (1993). A Response to Orzack and Sober: Formal Analysis and the Fluidity of Science. *Quarterly Review of Biology*. <https://doi.org/10.2307/3037251>
- Levy, A. (2018). Idealization and abstraction: Refining the distinction. *Synthese*, 13(1), 1–18. <https://doi.org/10.1007/s11229-018-1721-z>
- MacLeod, M., & Nagatsu, M. (2016). Model Coupling in Resource Economics: Conditions for Effective Interdisciplinary Collaboration. *Philosophy of Science*, 83(3), 412–433. <https://doi.org/10.1086/685745>
- Mäki, U. (2009). MISSing the world. Models as isolations and credible surrogate systems. *Erkenntnis*, 70(1), 29–43.
- Matthewson, J. (2011). Trade-offs in model-building: A more target-oriented approach. *Studies in History and Philosophy of Science Part A*, 42(2), 324–333.
- Matthewson, J., & Weisberg, M. (2009). The structure of tradeoffs in model building. *Synthese*, 170(1), 169–190. <https://doi.org/10.1007/s11229-008-9366-y>
- Mitchell, S. D. (2009). *Unsimple Truths*. University of Chicago Press.
- Muldoon, R., Smith, T., & Weisberg, M. (2012). Segregation that no one seeks. *Philosophy of Science*, 79(1), 38–62.
- Morgan, M. S., & Morisson, M. (1999). *Models as Mediators*. Cambridge University Press.
- Nguyen, J. (2020). It's Not a Game: Accurate Representation with Toy Models. *The British Journal for the Philosophy of Science*, 71(3), 1013–1041. <https://doi.org/10.1093/bjps/axz010>
- Odenbaugh, J. (2003). Complex Systems, Trade-Offs, and Theoretical Population Biology: Richard Levins's "Strategy of Model Building in Population Biology" Revisited. *Philosophy of Science (proceedings)*, 70(5), 1496–1507. <https://doi.org/10.1086/377425>
- Orzack, S. H., & Sober, E. (1993). *A Critical Assessment of Levins's The Strategy of Model Building in Population Biology (1966)*. 68(4), 533–546.
- Peters, R.H. (1991). *A Critique for ecology*. Cambridge University Press.
- Phillips, R. P., Ibáñez, I., D'Orangeville, L., Hanson, P. J., Ryan, M. G., & McDowell, N. G. (2016). A belowground perspective on the drought sensitivity of forests: Towards improved understanding and simulation. *Forest Ecology and Management*, 380, 309–320.
- Puccia, C. J., & Levins, R. (1986). *Qualitative Modeling of Complex Systems*. Harvard University Press.
- Railsback, S. F., & Grimm, V. (2011). *Agent-based and individual-based modeling: A practical introduction*. Princeton University Press.
- Reed, L. J., & Pearl, R. (1927). On the summation of logistic curves. *Journal of the Royal Statistical Society*, 90(4), 729–746.
- Reutlinger, A., Hangleiter, D., & Hartmann, S. (2018). Understanding (With) Toy Models. *The British Journal for the Philosophy of Science*, 69, 1069–1099. <https://doi.org/10.1093/bjps/axx005>

- Rosen, G. (2017). Abstract objects. In E. N., Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2017 Edition). <https://plato.stanford.edu/archives/spr2017/entries/abstract-objects/>
- Rothstein, R. (2017). *The color of law: A forgotten history of how our government segregated America*. Liveright Publishing.
- Sampson, R. J. (2019). Neighbourhood effects and beyond: Explaining the paradoxes of inequality in the changing American metropolis. *Urban Studies*, 56(1), 3–32.
- Scheiner, S. M., & Willig, M. R. (2008). A general theory of ecology. *Theoretical Ecology*, 1(1), 21–28.
- Schelling, T. C. (1969). Models of segregation. *The American Economic Review*, 59(2), 488–493.
- Shrader-Frechette, K. S., & McCoy, E. D. (1993). *Ecological concepts are problematic*. In *Method in Ecology: Strategies for Conservation* (pp. 11–67). Cambridge University Press.
- Sobek-Swant, S., Kluzza, D. A., Cuddington, K., & Lyons, D. B. (2012). Potential distribution of emerald ash borer: What can we learn from ecological niche models using Maxent and GARP? *Forest Ecology and Management*, 281, 23–31. <https://doi.org/10.1016/j.foreco.2012.06.017>
- Stillman, R. A., Railsback, S. F., Giske, J., Berger, U., & Grimm, V. (2015). Making Predictions in a Changing World: The Benefits of Individual-Based Ecology. *BioScience*, 65(2), 140–150. <https://doi.org/10.1093/biosci/biu192>
- Stockwell, D. (1999). The GARP modelling system: Problems and solutions to automated spatial prediction. *International Journal of Geographical Information Science*, 13(2), 143–158. <https://doi.org/10.1080/136588199241391>
- Strevens, M. (2008). *Depth*. Harvard University Press.
- Suárez, M. (2010). Scientific representation. *Philosophy. Compass*, 5(1), 91–101.
- Tompkins, D. M., & Veltman, C. J. (2006). Unexpected consequences of vertebrate pest control: Predictions from a four-species community model. *Ecological Applications: A Publication of the Ecological Society of America*, 16(3), 1050–1061.
- Toon, A. (2012). *Models and Representation* (pp. 6–33). Palgrave Macmillan UK.
- Travis, J., Coleman, F. C., & Auster, P. J. (2014). 'Integrating the invisible fabric of nature into fisheries management.' *Proceedings of the National Academy of Sciences of the United States of America*, 111, 2.
- Van Fraassen, B. C. (1980). *The Scientific Image*. Clarendon Press.
- Vinković, D., & Kirman, A. (2006). A physical analogue of the Schelling model. *Proceedings of the National Academy of Sciences*, 103(51), 19261–19265.
- Wang, L., & Jackson, D. A. (2014). Shaping up model transferability and generality of species distribution modeling for predicting invasions: Implications from a study on *Bythotrephes longimanus*. *Biological Invasions*, 16(10), 2079–2103. <https://doi.org/10.1007/s10530-014-0649-6>
- Ward, E. J., Holmes, E. E., Thorson, J. T., & Collen, B. (2014). Complexity is costly: A meta-analysis of parametric and non-parametric methods for short-term population forecasting. *Oikos*, 123(6), 652–661.
- Weisberg, M. (2007). Three kinds of idealization. *The Journal of Philosophy*, 104(12), 639–659
- Weisberg, M. (2013). *Simulation and Similarity*. OUP USA.
- Ylikoski, P., & Aydinonat, N. E. (2014). Understanding with theoretical models. *Journal of Economic Methodology*, 21(1), 19–36. <https://doi.org/10.1080/1350178X.2014.886470>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.