

Chapter 4

Modelling in Population Biology

4.1 Overview

While the previous chapter focused on lofty theoretical concerns regarding a particular variety of computational modelling methodology, the current chapter turns toward the pragmatic concerns facing simulation researchers in the discipline of population biology. Mathematical modelling has been popular for some time with population biologists, providing as it does a relatively simple means for tracking trends in natural populations, but the recent advent of agent-based modelling has added new possibilities for those in the community who seek a more detailed picture of natural populations than mathematical models can provide.

Population biologists long had doubts about the general usefulness of simulation, as evidenced by Richard Levins' criticisms of mathematical modelling techniques (Levins 1966). Levins argued that such methods faced fundamental limitations which would prevent simulation researchers from producing usable detailed models of populations. In the years following, debate has continued over Levins' ideas, particularly given the possible application of those same criticisms to more sophisticated computational modelling methods (Orzack and Sober 1993; Odenbaugh 2003).

After elucidating the most vital points raised during the lengthy Levins debate, this chapter will investigate how these arguments pertain to methodologies such as Alife and agent-based modelling. While Levins did originally focus upon mathematical modelling, the pragmatic concerns he raises regarding modelling as an enterprise remain useful when attempting to construct detailed, yet tractable, simulation models of natural phenomena.

Figures in this chapter were provided by Prof Seth Bullock, and are reprinted with permission.

These concerns will be relevant as well during the remaining chapters, in which we will begin the Part II of this text and delve into the issues surrounding social simulation. Levins' pragmatic concerns will allow us to draw a contrast between the issues of agent-based modelling in Alife and biology with the concerns of modellers within the social sciences. The framework described here will give us valuable background for these comparisons, demonstrating methodological issues that are important both to empirical research and to agent-based modelling. This framework will also be useful to keep in mind in Part III, given the methodological similarities between population biology and demography.

4.2 Levins' Framework: Precision, Generality, and Realism

4.2.1 Description of Levins' Three Dimensions

Levins' 1966 paper 'The Strategy of Model Building in Population Biology' (Levins 1966) triggered a long-running debate within the biological and philosophical community regarding the difficulties of successful mathematical modelling. Levins argued that during the construction of such models, population biologists must face the challenging task of balancing three factors present in any model of a natural system: precision, generality, and realism. He posits that only two of these three dimensions of a given model can be present to a significant degree in that model; strengthening two of these dimensions comes only at the expense of the third.

For example, to return to our bird migration example, a very precise and realistic model of the migration habits of a particular species of bird would necessarily lose generality, as such a model would be difficult to apply to other species. Likewise a broad-stroke model of migration habits would lose precision, as this model would not be able to cope with the many possible variations in migratory behaviour between species. Thus, a modeler must face the prospect of sacrificing one of these three dimensions when producing a model of a biological system, as no model could successfully integrate all three into a cohesive whole.

In order to demonstrate more clearly the contrasting nature of these three modelling dimensions, Levins outlines a hierarchy of models which biologists can construct within this framework (Table 4.1). First, Type I models (referred to hereafter as L1) sacrifice generality for precision and realism; these models may be useful in situations in which the short-term behaviour of specific populations is relevant to the problem at hand. Second, Type II models (L2) sacrifice realism

Table 4.1 Summary of Levins' three modelling types

Levinsian modelling framework	
L1	Precision and realism at the expense of generality
L2	Generality and precision at the expense of realism
L3	Generality and realism at the expense of precision

for generality and precision; these models could be useful in ecological models, for example, in which even heavily idealised models may still produce interesting results. Finally, Type III models (L3) sacrifice precision for generality and realism; models of this sort could be useful as general population biology models, in which realistic behaviour at the aggregate level is more important than accurate representation of individual-level behaviour. In a follow-up to his original description of these three dimensions, Levins clarifies his position that no model could be equally general, realistic and precise simultaneously, saying that a model of that sort would require such a complex and interdependent system of linked equations that it would be nearly impossible to analyse even if it did produce reasonable results (Levins 1968).

At the time of its publication, such a lucid description of the practicalities of model-building within population biology was unusual. Wimsatt (2001) lauds Levins for the way in which he 'talked about strategies of model-building, which philosophers never discussed because that was from the context of (presumably non-rational) discovery, rather than the (seemingly endlessly rationalise-able) context of justification' (p. 103). Levins (1968) book sought to further refine his ideas, though since then philosophers within biology and related disciplines have begun to take his ideas on board.

Of course Levins did produce this argument as a specific criticism of mathematical modelling efforts, but his description of these three key modelling dimensions is general enough in scope to provide a wealth of interesting implications for modern computational modelling. In an era when computer simulation is fast becoming the methodology of choice across numerous disciplines, such a cohesive argument outlining the fundamental limitations of modelling natural systems remains vital, and the strong criticisms of this framework coming from various corners of biology and philosophy serve an equally important role in bringing Levins' ideas up-to-date with the current state-of-the-art.

4.3 Levins' L1, L2 and L3 Models: Examples and Analysis

4.3.1 *L1 Models: Sacrificing Generality*

As in the variation of our bird migration example presented above, studies which focus upon realism and precision must necessarily decrease their generality under Levins' framework. Given that our hypothetical realistic and precise bird-migration model aims only to examine the migration habits of a particular bird species, the modeler will have greater difficulty in applying his results to other species. This is not to say that his data may not also explain by coincidence the migration habits of other bird species, but rather that his model displays those habits using the specific data of one species and does not set out with the intent of describing migration in a broader context.

4.3.2 L2 Models: Sacrificing Realism

An L2 model under Levins' framework emphasizes generality and precision at the expense of realism. In this case, the bird modeler would reduce his connection to real data, hoping to produce a model which, while displaying deviations from empirical data regarding specific species or short-term migration situations, can provide a useful set of parameters or equations that can describe the phenomenon of bird migration in an idealised context.

Levins notes that physicists entering the field of population biology often operate in this way, hoping that their idealised models may stray closer than expected to the equations that govern the overall behaviour they are investigating (Levins 1966). While this may seem misguided initially, a successful validation of data from such a model using empirically-gathered data from the real system could certainly provide a valid means for revising the model in a more realistic direction. In this way we might say that the physicist is tacitly accepting the fundamental limitation of such an idealised model by making no attempt to attach realism to his formalisation of the problem, and instead seeks to reach a more realistic perspective on the problem by integrating insights produced from 'real' data.

4.3.3 L3 Models: Sacrificing Precision

Taking our bird example in a slightly different direction, imagine that the population biologist chooses to eschew precision in favour of generality and realism. In this case, he is not concerned with matching his model's data with that of the real-world system; instead he is interested in more simple comparisons, describing the behaviour of the birds in response to changes in environmental stimuli for example. The simulation may state that the birds migrate more quickly during particularly cold winters, for example, which allows him to draw a general conclusion about the behaviour of real birds during such times; the actual number of birds migrating, or the manner in which they migrate, is unimportant for drawing such conclusions in this context. Levins views this method as the most fruitful for his field, and accordingly follows this methodology himself.

4.4 Orzack and Sober's Rebuttal

4.4.1 The Fallacy of Clearly Delineated Model Dimensions

Orzack and Sober's rebuttal to Levins' proposals begins by attempting to clarify Levins' descriptions of generality, precision and realism (Orzack and Sober 1993).

In their view: generality refers to a greater correspondence to real-world systems; realism refers to the model's ability to take account of more independent variables; and precision refers to the model's ability to make valid predictions from the relevant output parameters. Further, Orzack and Sober contend that a model's parameters can be specified or unspecified, specified models being those in which all parameters are given specific values. Importantly they consider unspecified models as necessarily more general, given the lack of assigned parameter values. Orzack and Sober argue that these characteristics of models interact in unexpected ways, and the resulting deviations from Levins' framework may provide a means for modelers to escape that framework.

4.4.2 Special Cases: The Inseparability of Levins' Three Factors

With these ideas in mind, Orzack and Sober propose that if any model is a special case of another, then the special case must necessarily start from the same position amongst Levins' three dimensions as the original model (at least according to their definitions of generality, realism and precision). The new model thus will gain in one or more of the Levinsian dimensions without losing in the others (which would be the case under Levins' formulation). This obviously clashes with Levins' thesis, and as a result Orzack and Sober claim that, in some cases, these three properties are not connected, meaning that a trade-off between them is not necessary.

In this way, assuming Orzack and Sober's criticisms hold true, the modeler could continually refine a generic modelling framework related to a certain problem to evade Levins' stated problems. For example, if a general unspecified model of our bird-migration problem was refined by taking into account new biological data, thus producing potential fixed values for the parameters of the model, then this specified version of that model has increased its realism without losing the generality and precision of its more generalised parent. Orzack and Sober would argue that the model has necessarily started from the same level of correspondence to real-world systems, ability to account for independent variables, and ability to make predictions as its parent, and as a special case of that model has gained in realism without sacrificing the other two dimensions in the process.

Despite Orzack and Sober's vehement position that Levins' framework of modelling dimensions is fundamentally flawed, their criticisms do seem unsatisfactory. They neglect much mention of Levins' points regarding the practicalities of model-building, and how these three dimensions can influence that process, and his related points regarding the inherent difficulty involved in analysing overly-complex models.

4.5 Resolving the Debate: Intractability as the Fourth Factor

4.5.1 *Missing the Point? Levins' Framework as Pragmatic Guideline*

Thus, despite the validity of Orzack and Sober's concerns regarding Levins' three dimensions as hard constraints on the modelling process, their complete deconstruction of his hypotheses in rigorous semantic terms may be excessive. Levins' points regarding the practicalities of model-building seem vital to an understanding of his underlying convictions; his framework seems to provide a summation of what, in his view, influences the usefulness and applicability of a given model to a particular problem.

4.5.2 *Odenbaugh's Defence of Levins*

Odenbaugh (2003) defends Levins against the semantic arguments of Orzack and Sober by first noting the general peculiarities of Levins' writing style which may lead the reader to draw more forceful conclusions than those originally intended. As a Marxist, Levins' mention of 'contradictory desiderata' in relation to the tension between these three modelling dimensions is not used to imply that these properties are mutually exclusive or entirely logically inconsistent, but rather that these model properties can potentially influence one another negatively. Such terminology is common amongst those of a Marxist bent, but other readers may assume that Levins' description of contradictory model properties does imply a hard mutual exclusivity.

Odenbaugh also concedes that theoretically one may find models in which trade-offs between generality, realism and precision are unnecessary, or perhaps even impossible. However, he argues that Levins' ideas are intended as pragmatic modelling guidelines rather than definitive constraints as implied by Orzack and Sober; Levins concerns himself not with formal properties of models but instead with concepts that may guide our construction of sensible models for natural systems. Odenbaugh concludes his analysis of Levins' thesis by summarising his view of Levins' intended purpose:

...Levins' discussion of tradeoffs in biological modelling concerns the tensions between our own limitations with respect to what we can compute, measure and understand, the aims we bring to our science, and the complexity of the systems themselves. (Odenbaugh 2003, p. 17)

Thus in one sense Levins' ideas stretch beyond the formal guidelines Orzack and Sober accuse him of constructing, but instead encompass statements regarding the fundamental limits of tractability and understanding which underlie all modelling endeavours.

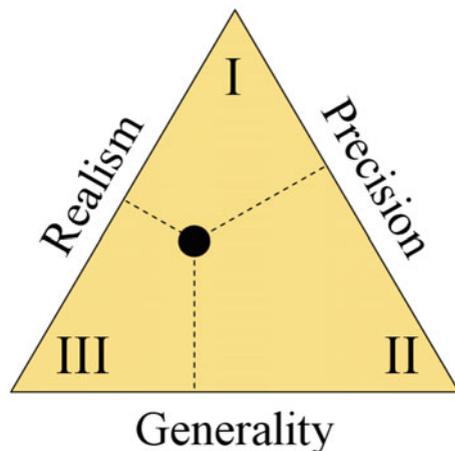
4.5.3 *Intractability as the Fourth Factor: A Refinement*

Odenbaugh's defence of Levins gives us an important additional concept that may be useful for refining Levins' modelling framework. He mentions the importance of computability, and the tensions that such practical concerns can produce for the modeller. Of course, as already intimated by Levins, a model's amenability to analysis is also of primary concern for the modeller in any field. Odenbaugh makes a valuable point, demonstrating that Levins' modelling dimensions are of primary importance to the modeller seeking pragmatic guidance in the task of constructing a model that can be successfully analysed. Here we use the term 'tractability' to refer to this characteristic of a given model; while other terms such as 'opacity' may also be appropriate, 'tractability' carries the additional connotation of mathematical or computational difficulty, and thus seems preferable in the context of Odenbaugh's formulation.

While Odenbaugh's synthesis of Levins' claims does greatly diminish the effect of Orzack and Sober's analysis on the validity of Levins' claims, one might take this rebuttal even further, pointing to Levins' ideas regarding tractability as a potential solution to this semantic dilemma. While the hypothetical models proposed by Orzack and Sober are certainly logically possible, they ignore one of Levins' central concerns: if one fails to make a trade-off between these three dimensions while constructing a model, the model will become extremely complex and thus difficult, or impossible, to analyse. In a sense, tractability becomes a fourth dimension in Levins' formulation: the closer one gets to a complete balance of generality, precision and realism, the further one moves towards intractability.

To clarify this relationship somewhat, we might imagine Levins' three model dimensions as forming the vertices of a triangle, with each vertex representing one of the three dimensions (Fig. 4.1). When we construct a model, that model occupies some place within that triangle, its position indicating its relative focus on each of

Fig. 4.1 Levins' three model dimensions



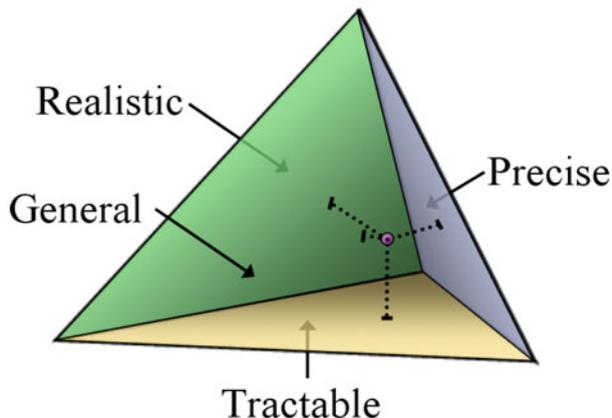


Fig. 4.2 Levins' modelling dimensions with the addition of tractability

those three dimensions. In this arrangement, moving a model toward one side of the triangle (representing a focus upon two of the three model dimensions) would necessarily move it away from the third vertex, representing the trade-off between generality, precision and realism.

However, as per this clarification which places tractability as a fourth dimension in Levins' modelling framework, we may add a fourth vertex to our triangle, extending it to a tetrahedron. In this arrangement a model which tends toward the centre of the triangle, representing an equal balance of all three initial dimensions, could be seen as approaching that fourth dimension of intractability (Fig. 4.2). The placement of intractability into the modelling dimensions is critical, as this allows for models which balance all three Levinsian factors to exist; the cost of doing so becomes a loss of tractability rather than an outright declaration of impossibility.

Of course, any diagram of such an interplay between such broad concepts will fall short of demonstrating the full import of the ideas described here, not least because such a diagram requires an illustration of four interacting dimensions. This diagram does not entirely avoid such difficulties, as one can clearly imagine paths within such a space which avoid the pitfalls described by Levins. However, this representation proves preferable to other potential approaches due to its useful illustration of the space of models described by Levins' dimensions. If we imagine that any movements of the model in question toward the balanced centre of the Levinsian triangle result in a related movement toward tractability and the upper vertex of our tetrahedron, then the representation becomes more clear.

In any case, this diagram of the delicate balance of all four factors seeks to express Levins' idea, as clarified by Odenbaugh (2003), that his pragmatic concerns represent a fundamental limitation on the part of the researcher rather than the model

itself. While a model may potentially encompass all three dimensions equally, the parameters of the model will become intractable and impossible to analyse to the point of exceeding the cognitive limitations of the scientist. The exponentially more difficult task of analysing the model in such a circumstance eliminates the time-saving benefits of creating a model in the first place, necessitating a carefully-considered balancing of the computational demands of a model with its consistency along these dimensions of realism, generality and precision.

4.6 A Levinsian Framework for Alife

4.6.1 Population Biology vs. Alife: A Lack of Data

While the pragmatic concerns illuminated by Levins are clearly relevant for the biological modeller, there are significant differences between the mathematical modelling approach and the Alife approach as discussed in the previous chapter. Mathematical models for population biology often begin with field-collected data that is used to derive appropriate equations for the model in question. For example, our bird migration researchers may wish to establish a model of migration rates in a bird population using these more traditional methods, rather than the computational examples of possible bird migration studies examined in the previous chapter. In this case, they would first perform field studies, establish patterns of behaviour for those populations under study, and tweak their model parameters based upon that initial empirical data.

By contrast, many Alife studies may begin with no such background data. For example, a modeller that wishes to examine the development of signaling within animal populations would have great difficulty obtaining empirical data relevant to such a broad question. Looking at such an extensive evolutionary process is naturally going to involve a certain scarcity of usable empirical data; after all, watching a species evolve over thousands of generations is not always a practical approach. Thus the researcher must construct a simulation using a set of assumptions which best fit current theoretical frameworks about the development of signalling. Given the methodology of agent-based simulation in particular, many of these simulations will be programmed using a ‘bottom-up’ approach in which low-level factors are modeled in the hopes of producing high-level complexity from the interactions of those factors.

Of course, in order to produce results in a reasonable period of time which are amenable to analysis, such models must be quite simplified, often making use of highly-idealised artificial worlds with similarly simplified rules governing those organisms within the simulation. With these simulations producing reams of data relating only to a highly-simplified version of real-world phenomena, will these

models necessarily obey Levins' thesis regarding the trade-off between generality, precision and realism? For example, need an Alife modeler concern himself with such factors as 'precision' when the model cannot be, and is not intended to be, a realistic representation of a specific animal population?

4.6.2 Levinsian Alife: A Framework for Artificial Data?

Despite the inherent artificiality of data produced within Alife, Levins' framework can still apply to such a methodology. An Alife model designed in a highly abstract fashion, as in the above communication example, can simply be placed amongst the L3 models, as it aims to produce data which illuminates general trends in organisms which seek to communicate, rather than orienting itself towards a particular population. Some models may also use a specific species as a basis for an Alife simulation, thus leading that model away from generality and toward realism and precision. Similarly, overly complex simulations, like overly complex mathematical models, become difficult to compute and hence cumbersome; this implicitly limits the precision of an artificial life simulation.

However, as hinted above and as described by Orzack and Sober (1993), the simulation researcher will run into trouble regarding the meaning of these three dimensions within this trade-off. One might question the utility of regarding any simulation which deals with broad-stroke questions using idealised worlds and populations as 'realistic' or 'precise' in any conventional sense. Despite this problem, as Taylor describes, Levins did provide a caveat that 'models should be seen as necessarily 'false, incomplete, [and] inadequate' but 'productive of qualitative and general insights' (Taylor 2000, p. 197). With this in mind we might take Levins' framework as a very loose pragmatic doctrine when applied to computer simulation; we might ask whether a simulation resembles reality rather than accurately represents reality.

4.6.3 Resembling Reality and Sites of Sociality

This pragmatic application of the general thrust of Levins' framework might seem a relief to simulation researchers and theorists, but still a great problem remains. Determining whether a simulation resembles reality is far from straightforward, and might be evaluated in a number of different ways. The evaluation of a model as a reasonable resemblance to the real system which inspired it could appear quite different depending on the design of the simulation itself or the appearance of the data produced by the simulation, assuming of course that any conventional statistical data is produced at all. This can easily result in rather overdramatic or otherwise inappropriate conclusions being drawn from the study in question.

Take, for example, two groups of researchers, each attempting a realistic and enlightening model of our migrating birds. Group A models the birds' physical attributes precisely using three-dimensional computer graphics, models their movements in exacting detail using motion-capture equipment on real birds, and translates it all into a breathtaking display of flying birds on the monitor. Meanwhile, Group B models the birds only as abstract entities, mere collections of data which move, migrate and reproduce only in a computational space created by skillful programming. The model uses pre-existing data on bird migrations to produce new insights about why and how birds tend to migrate throughout the world.

In this example, how can we say which model is more realistic or has a more pronounced resemblance to reality? Some may claim that Group A's model provides a provocative and insightful look into how birds move through the air, while others may claim that group B's model produces more exciting data for biologists and zoologists despite the abstractions applied to the birds within that simulation. In reality such a comparison is virtually meaningless without such a context, as each model is designed for an entirely different purpose. Each model may resemble certain elements of bird behaviour in certain enlightening ways, but neither takes into account the effects of factors addressed in the other model.

Taylor (2000) continues upon this line of thinking, noting the existence of varying 'sites of sociality' in which modelers must operate. These sites correspond to points at which social considerations within a scientific discipline begin to define parameters of the model, rather than considerations brought about by the subject matter of the model itself. Thus, if a zoology conference views Group A's graceful three-dimensional birds, they are likely to receive acclaim for their models accuracy. Similarly, if Group B presents their abstract migration model to a conference of population biologists, they're more likely to receive a warm reception for their work. In either case, if Group A and B switched places, those communities would likely make their presentations rather more difficult.

Such discussions are very relevant for many researchers within the computational modelling community. For example, Vickerstaff and Di Paolo's model of path integration in ants provides one instance of a model crossing between research communities (Vickerstaff and Di Paolo 2005). The model is entirely computational in nature, and thus provides no hard empirical data, and yet the model was accepted and published by the *Journal of Experimental Biology*. The authors describe the process of acceptance as an interesting one; the editors of the journal needed to be convinced that the model was relevant enough and the ideas it presented novel enough to warrant the interest of the journal's readership.

Taylor's discussion of social considerations within given scientific disciplines makes this an interesting happening. If the model had been pitched slightly differently, say by using much more complex neural models, or making vastly inappropriate biological assumptions in the model's implementation, then this group would likely have been less receptive regarding the model's publication for their readership. For such a community, biological relevance is highly important; consider the contrast with journals like *Artificial Life*, for example, in which the presentation of computational models is far more broad in nature.

4.6.4 Theory-Dependence Revisited

This idea of sites of sociality harkens back to the previous chapter's discussion of theory-dependence in science. The construction of any model, whether mathematical, computational, or conceptual, is an inherently theory-dependent exercise. All varieties of models must simplify elements of reality in order to study it in these reduced forms; otherwise the modeler has gained little from avoiding the more costly and time-consuming methods of empirical data-gathering which may produce similar results.

With all of these abstractions and simplifications, however, a certain resemblance to reality must be maintained. Our detailed three-dimensional model of flying birds maintains a useful resemblance to real-world flying behaviours; likewise, our model of migrating birds, while more abstracted, may retain that useful resemblance by demonstrating migration patterns observed in real bird populations, or even demonstrating the functioning of this higher-level behaviour amongst a generalised population (as is the goal in much of *Alife*).

Determining whether a model resembles reality is thus an inherently pragmatic process, depending upon both the structure of the model itself and the structure of its intended audience. Some questionable elements of a model may become obvious when the model is presented, but on the other hand subtle assumptions that may affect the model's results will be far less obvious. This makes evaluation of that model even more difficult, as extracting the nature of a model's theory-dependence is not always straightforward when observing the results, and even then different audiences may perceive those theory-dependent elements differently. Of course this does not imply complete relativism in scientific results, but this view does stress the importance of the audience for those results in determining the overall success of a modelling effort.

4.7 Tractability Revisited

4.7.1 Tractability and Braitenberg's Law

As we have seen, our revised four-factor version of Levins' framework gives us a pragmatic theoretical background for building computer simulations in which all such simulations are fundamentally constrained by problems of tractability. This can lead to such situations as those described above, in which greatly-simplified models are far more tractable but also in turn much more abstracted, making for potentially provocative and misleading results. However, as indicated in our examinations of *Alife* simulations, there does seem to be a greater degree of such grandiose approaches in computer simulation than in mathematical modelling. What about computer simulation leads some to attempt to achieve provocative results at the expense of tractability?

The answer may lie in a fundamental difference in the construction of computer simulations and mathematical models. Braitenberg's influential work *Vehicles* (Braitenberg 1984) provides an insight into such elements of model construction. In this work, he outlines a series of thought experiments in which he describes the construction of highly-simplified vehicles, each incorporating only basic sensors, which can nevertheless display surprisingly complex behaviour that is startlingly reminiscent of living creatures. The most famous example of these is his light-seeking vehicle, which is drawn towards a light-source simply by having a light-sensor directly hooked to its wheels, producing a simple light-seeking behaviour reminiscent of that of moths.

During the course of the book Braitenberg proposes an idea known as the 'law of uphill analysis and downhill invention,' describing the inherent difficulty in capturing complex behaviour from an external perspective:

... it is much more difficult to start from the outside and try to guess internal structure just from the observation of behaviour.... It is pleasurable and easy to create little machines that do certain tricks. [But] analysis is more difficult than invention just as induction takes more time than deduction.

For Braitenberg then, his vehicular 'experiments in synthetic psychology' may have provided an easier path to insight than dutiful, and most likely tedious, observation and inference from the behaviour in a real system.

Braitenberg's Law illuminates the difference between the factors of analysis and invention in mathematical models and computer simulations. When constructing a mathematical model, this construction (invention) is quite tightly coupled to the analysis that will result from the model, as a mathematical model which is nearly impossible to analyse will be correspondingly difficult to modify and tweak. As Braitenberg describes, the 'invention' is easier than the analysis, and in the case of a mathematical model the difficulty in analysis will also affect that invention process.

When constructing a computer simulation the disconnect between these two factors is much larger. In the case of agent-based models, each agent is constructed as a confluence of various combined assumptions laid out during the 'invention' process, and once the simulation begins those agents will interact in non-trivial ways both with other agents and with the related virtual environment. This results in a troublesome opacity for the analyst, as the mechanics and the results of the simulation do not necessarily have a direct correspondence.

For example, say our abstracted bird migration model from earlier used neural networks to drive the simulated birds' movements and their responses to environmental stimuli. When analysing the results, determining how those complex sets of neural connection weights correspond to the observed higher-level behaviours is a herculean task, and even more so when we consider how those weights change in response to a complex and changeable virtual environment. Even though these network weights influence the behaviour of each virtual bird, and thus have a direct impact upon that simulated environment, the relationships of those weights to the simulation results are far from obvious.

Clearly then the coupling between the synthesis and analysis of an agent-based model is significantly looser than that of a mathematical model. As these simulations continue to grow in size and complexity, the resultant opacity of these models could become so severe that synthesis actually forgoes analysis. In this nightmare scenario the process of analysing such an opaque simulation can become so time-consuming that the time saved by running that simulation is eclipsed by the time wasted when trying to penetrate its opacity.

4.7.2 *David Marr's Classical Cascade*

A useful perspective on this tractability problem comes from artificial intelligence via David Marr's description of hierarchical levels of description (Marr 1972). Marr discusses three levels at which the researcher develops explanations of behaviour: level 1 (hereafter referred to as M1), in which the behaviour of interest is identified computationally; level 2 (M2), in which an algorithm is designed that is capable of solving this behavioural problem; and level 3 (M3), in which this behavioural algorithm is implemented (Table 4.2). The agent-based simulation methodology tends to deviate from this hierarchy, however, as many such models are not so easily described in terms of a single algorithm designed to solve a single behavioural problem.

Peacocke's extended version of Marr's hierarchy seems more applicable to our concerns (Peacocke 1986). He adds a level between M1 and M2 in Marr's hierarchy, appropriately termed level 1.5 (M1.5), in which the modeler adds to his initial M1 specification of his problem in computational terms by drawing upon a 'body of information' appropriate to the behaviour of interest. For example, our bird migration example from earlier may incorporate an artificial neural network which drives the behaviour of individual agents within the model, in addition to the other aspects of the agents' environment that drives their behaviour. In this case we define M1.5 to include both the original formulation of the model and the resources the model draws upon to solve the behavioural problem of interest, which for our bird migration model would be the artificial neural network driving each agent.

From constructing our bird migration model, we would naturally proceed to running the simulation and obtaining results for analysis. However, under Peacocke's analysis, we have essentially just skipped the M2 portion of Marr's hierarchy entirely. We specified our problem computationally, constructed a model using an appropriate body of information, and then proceeded directly to the implementation

Table 4.2 Summary of Marr's levels

Marr's levels of description	
M1	Problem identified computationally
M2	Algorithm designed to solve the problem
M3	Behavioural algorithm is implemented

step without developing a distinct algorithm. Unfortunately, M2 is identified as the level of this hierarchy which produces the most insight into the behaviour of interest; that M2 algorithmic understanding of the problem would normally allow the researcher to develop a more in-depth explanation of the resultant behaviour.

4.7.3 Recovering Algorithmic Understanding

With this problem in mind, agent-based modelers may choose to attempt to step backwards through the classical cascade, proceeding from M3 back to M2 in order to achieve this useful algorithmic understanding that is denied to them by this simulation methodology. However, as noted earlier, these models already sit in a difficult position between synthesis and analysis. How might we jump backwards to divine this algorithm, when our model may already be too complex in its synthesis to promote the type of analysis which may produce that algorithmic understanding?

Perhaps a better course would be to ask why an agent-based modeler needs this type of explanation at all. Agent-based models are most often developed to produce emergent behaviour, which by its very nature is not particularly amenable to algorithmic reduction. A complex high-level behaviour deriving from simple low-level interactions is not something easily quantified into a discrete form.

Similarly, the rules which define a given simulation may not lend themselves easily to such analysis either; the results of a simulation may be too opaque to narrow down the specific influences and effects of any of the low-level rules which produce that emergent behaviour. Alternatively, if the researcher produces a model which bears a useful resemblance to the behaviour of interest as seen in the natural system, then this validation against empirical results may be enough of a confirmation of the model's concept that these complicated analyses may not be viewed with much importance.

4.7.4 Randall Beer and Recovering Algorithmic Understanding

Of course, while recovering this algorithmic understanding may be difficult for the agent-based modeller, this is by no means an impossible task. For example, even the simplest artificial neural network can be quite opaque when attempting a detailed analysis of the behaviour of that network; neural network connection weights can provide novel solutions to certain tasks within a model, but unraveling the meaning of the connection weights in relation to the problem at hand can take a great deal of effort and determination.

Despite these difficulties, such analyses have been performed, but only after a great deal of concentrated effort using novel cluster-analysis techniques and similar methods. A seminal example of this is Randall Beer's examination of a minimally cognitive agent performing a categorical perception task (Beer 1995, 2003b). Beer's

agent employs a network of seven neurons to discriminate between falling objects, tasked with catching circular objects while avoiding diamond-shaped ones. The best evolved agents were able to discriminate between objects 99.83% of the time after 80 generations of evolution. Beer's goal was to develop a minimally-cognitive agent performing a vital, but simple, task, then analyse the simulation to understand how the task is successfully performed, and the psychophysical properties of the agent.

In order to study the agent and its performance as part of this coupled brain-body-environment system, Beer analyses the agent and its behaviour in excruciating detail using the language of dynamical systems theory. He argues that such analysis must form the backbone of a proper understanding of any given simulated phenomenon; as Beer states, 'it is only by working through the details of tractable, concrete examples of cognitive phenomena of interest that we can begin to glimpse whatever general theoretical principles may exist' (Beer (2003b), p. 31, also see Beer 2003a).

Beer's analysis is not without its detractors, even in this case. Webb provides a lengthy treatise on the lack of biological relevance in Alife models, using Beer as a central example (Webb 2009). He contends that while Beer's analysis is amply justified in his paper and the relevant responses (Beer 2003a,b), he contradicts himself when discussing the issue of the relevance of his model. He notes Beer's contention that he does not seek to create a serious model of categorical perception, and yet this contradicts his later statements regarding the value of his model in discovering properties of categorical perception; thus, Webb says, 'empirical claims are being made about the about the world based on the model results' (Webb 2009, p. 11). In Webb's view, Beer attempts to take his analysis a step beyond where he himself claims it has actual validity. He does not wish his simulation to be considered a true model of a natural system, and yet wishes also to use his analysis to make claims as if it is such a model.

Webb's discussion is an important one, as it bears upon our further discussion of the difficulties of artificial worlds. Beer's analysis is detailed, and he largely succeeds in avoiding a major pitfall of models of this type and their lack of algorithmic understanding. Yet, in the process of doing so, he also places emphasis for the reader on another concern, namely the relevance of such models outside of an exercise in formal analysis or the study of a simulated system for its own sake.

4.7.5 The Lure of Artificial Worlds

This discussion of algorithmic understanding, or the lack thereof, in agent-based simulations provides an insight into the lure of these artificial worlds for the research community. By constructing these models the researcher has essentially replaced the nuances of the real system of interest with a simplified artificial version, avoiding the complexities and occasional impossibilities (both practical and financial) of attempting certain empirical studies. In addition, as noted earlier and as discussed by Webb, the study of this artificial world becomes a sort of pseudo-empirical activity on its own as the researcher must tweak parameter settings and observe many runs

of this compressed and abstracted version of the natural system. Of course these models can exhibit the aforementioned opacity to analysis, but many leave the black-box internals of their simulations alone, content to skip what is normally an essential step in the modelling of behaviours, and further a step acknowledged by many to be essential to developing an understanding of the behaviour in evidence, as noted above.

This tactic for avoiding the complex questions raised by a lack of algorithmic understanding in agent-based modelling is also helpful for those troubled by Levins' framework. For a relatively small initial investment for the researcher (small at least in comparison to large-scale empirical studies), one appears to gain a great deal of insight through this artificial world, and perhaps even avoid the issues of realism, generality and precision that limit the application of more traditional modelling methods. However, the fourth dimension of Levins' modelling framework adds a tractability ceiling to any modelling endeavour, regardless of scope or method; while an artfully-constructed model may allow an equal balance of the three basic Levinsian factors, this tractability ceiling will greatly limit the ability of that model to produce useful and analysable data.

Initially the use of artificial worlds in computational models appears to propel the model in question through the tractability ceiling. The modeler can integrate the three Levinsian factors much more easily in an artificial world, lacking as it does the complexities apparent in the natural world. Marr, Peacocke and Clark remove this possibility as well, however, by exposing the inability of the artificial-world modeler to achieve an algorithmic understanding of the system he chooses to model. In that sense an artificial world model can strive to be little more than a proof-of-concept, an indication of possible directions for empirical research but of very little use in deriving actual useful conclusions about the behaviour of interest.

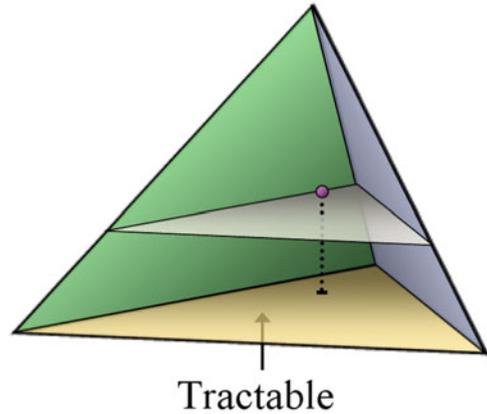
4.8 Saving Simulation: Finding a Place for Artificial Worlds

4.8.1 Shifting the Tractability Ceiling

Thus far our attempts to find a justification for models utilising artificial worlds have been rebuffed quite frequently; even the relatively promising framework for strong ALife outlined in Chap. 3 provides little insight into creating useful models of that type. Levins has criticised modellers for failing to balance generality, realism and precision in our models; our refinement of Levins' three factors has demonstrated the dangers of constructing a balanced, but intractable model; and Marr, Peacocke and Clark have pointed out the failing of certain models to develop an algorithmic understanding of a natural system or behaviour. How then might our artificial worlds of roving software agents contribute to the realm of empirical science?

Thankfully, agent-based simulations and similar methodologies do not exist in a vacuum. The tractability ceiling is a formidable obstacle, but not a static

Fig. 4.3 Breaking through the tractability ceiling



one: as computing power increases, for example, our ability to develop and run increasingly complex simulations becomes ever more feasible. In addition the continual refinement of theoretical frameworks in empirical science, which drive the simplifying assumptions that underlie the creation of these artificial worlds, can completely alter our understanding of a particular problem, and thus alter our methods for developing an understanding of that problem.

However, even with this shifting tractability ceiling (Fig. 4.3), we are left with substantial theoretical difficulties. How can we avoid the seemingly insurmountable problems raised by Marr, Peacocke and Clark? With simulations easily falling into the trap of opacity, how can we still claim that our models remain a useful method for empirical science?

4.8.2 *Simulation as Hypothesis-Testing*

Agent-based models may serve as a much more fertile method of research when used as a well-designed adjunct to more conventional empirical methods. Like Noble, Di Paolo and Bullock's 'opaque thought experiments,' simulations may function best as a method for testing the limitations of a hypothesis (Di Paolo et al. 2000). Our agent-based model of bird migration could provide a way to examine how existing theories of migration might predict the future movements of the real species, and through its results produce some potential avenues for empirical testing of those predictions.

The greatest problem lies, as we have seen, in using a simulation to produce useful empirical data on its own. While strong Alife theorists may choose to escape this problem by contending that their artificial worlds are in fact 'real' worlds full of their own digital biodiversity, more conventional approaches suffer under the constraints particular to the simulation methodology. Overly abstracted simulations

may not produce the desired resemblance to the natural systems, while overly complex simulations may actually preclude a sensible analysis of the resulting data.

Through all of these various arguments regarding the simulation methodology, the difficulty of empirical data-production remains the most prominent. Beyond the trap of opacity lies another trap: simulation models can all too easily lend themselves to being studied simply for their own sake. As these models brush against the tractability ceiling, they must necessarily fall short of the complexity of the real system of interest; the model cannot fundamentally replace the real system as a target for experimentation. Studying a model simply to test the boundaries and behaviours of the attached artificial world (as in Ray's *Tierra*, for example) may certainly lead to intriguing intellectual questions, but is unlikely to lead to substantive conclusions regarding the natural world.

4.9 Summary and Conclusion

Looking through the lens of population biology, we have seen the varied and disparate opinions dividing the simulation community. Levins' early debates about the merits and pitfalls of mathematical modelling continue today in a new form as computational modelling continues to grow in prominence throughout various disciplines.

The limitations outlined by Levins remain relevant precisely because they are linked to the fundamental shortcomings of any modelling endeavour, rather than being confined particularly to the realm of mathematical models. The three dimensions of generality, precision and realism provide a useful pragmatic framework under which the simulation researcher and theorist can examine the fundamental assumptions of a given model, regardless of the paradigm under which it was constructed.

Our expansion of Levins' framework provides closure to this list of theoretical concerns, showing how the issue of tractability confines computational modelling. While simulations are a remarkably attractive methodology for their relative simplicity and apparent explanatory power, these characteristics cannot overcome the simple problems of complexity and analysability which can make an otherwise appealing model incomprehensible.

Even assuming tractability and the four dimensions of Levins' updated framework are examined in detail during the construction of a simulation, Marr, Peacocke and Clark point out the difficulties inherent in deriving useful insight from such a model. The usual procedure for developing an understanding of a natural system's behaviour is circumvented by the simulation process; the simulation produces results directly, without the intermediate step of producing an algorithmic understanding of the behaviour of interest, as is the normal case in empirical research. This leaves the simulation researcher in a very difficult position as he attempts to work backward from his completed simulation results to find that algorithmic understanding.

From this chain of investigations the overall appeal of the simulation methodology and the production of these ‘artificial worlds’ becomes apparent. By creating an artificial world in which simulated agents interact and produce complex behaviour of their own accord, the researcher can evade the questions raised by Levins and plausibly escape the constraints pointed out by Marr, Peacocke and Clark. The artificial world can easily become a target of experimentation, providing as it does a distilled and simplified version of its real-world inspiration, without some of the attendant analytical difficulties.

Of course, this methodology produces analytical difficulties of its own. This artificial world cannot replace the real world as a source of empirical data, as their complexity is far lower and their construction far more theory-dependent. As much as the simplicity and power of computer simulation exerts a powerful lure for the research community, the objections of Levins and others provide a powerful case for a great deal of caution in the deployment and use of such methods.

As looking through the lens of population biology provided a useful set of criticisms and possible constraints for simulation modelers, so looking through the lens of the social sciences will provide a look at some promising areas for simulation research. We have already seen how simulation can be very appealing for fields which suffer from a dearth of empirical data, and the social sciences are in the unique position of providing a great variety of theoretical background in a field which is often very difficult to examine empirically. An examination of the use of simulation in the social sciences will provide a means to integrate these varying perspectives on simulation into a framework that accounts for the strengths and weaknesses we have exposed thus far.

References

- Beer, R. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 172–215.
- Beer, R. (2003a). Arches and stones in cognitive architecture. *Adaptive Behavior*, 11(4), 299–305.
- Beer, R. (2003b). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4), 209–243.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In M. A. Bedau, et al. (Eds.), *Artificial life VII* (pp. 497–506). Cambridge, MA: MIT Press.
- Levins, R. (1966). The strategy of model-building in population biology. *American Scientist*, 54, 421–431.
- Levins, R. (1968). *Evolution in changing environments*. Princeton: Princeton University Press.
- Marr, D. (1972). Artificial intelligence – a personal view. In M. A. Boden (Ed.), *The philosophy of artificial intelligence* (pp. 133–147). Oxford: Oxford University Press.
- Odenbaugh, J. (2003). Complex systems, trade-offs, and mathematical models: A response to Sober and Orzack. *Philosophy of Science*, 70, 1496–1507.
- Orzack, S., & Sober, E. (1993). A critical assessment of Levins’ ‘The Strategy of Model Building (1966)’. *Quarterly Review of Biology*, 68, 534–546.

- Peacocke, C. (1986). Explanation in computational psychology: Language, perception and level 1.5. *Mind and Language*, 1, 101–123.
- Taylor, P. (2000). Socio-ecological webs and sites of sociality: Levins' strategy of model-building revisited. *Biology and Philosophy*, 15(2), 197–210.
- Vickerstaff, R. J., & Di Paolo, E. A. (2005). Evolving neural models of path integration. *Journal of Experimental Biology*, 208, 3349–3366.
- Webb, B. (2009). Animals versus animats: Or why not the real iguana? *Adaptive Behavior*, 17(4), 269–286. <https://doi.org/10.1177/1059712309339867>.
- Wimsatt, W. C. (2001). Richard Levins as philosophical revolutionary. *Biology and Philosophy*, 16, 103–108.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), 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 license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license 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.

