

# Chapter 11

## Combining Mathematical and Simulation Approaches to Understand the Dynamics of Computer Models

Luis R. Izquierdo, Segismundo S. Izquierdo, José M. Galán,  
and José I. Santos

**Why Read This Chapter?** To learn how to better understand the dynamics of computer models using both simulation and mathematical analysis. Our starting point is a computer model which is already implemented and ready to be run; the objective is to gain a thorough understanding of its dynamics. Combining computer simulation with mathematical analysis can help to provide a picture of the model dynamics that could not be drawn by using only one of the two techniques.

**Abstract** This chapter shows how computer simulation and mathematical analysis can be used together to understand the dynamics of computer models. For this purpose, we show that it is useful to see the computer model as a particular implementation of a formal model in a certain programming language. This formal model is the abstract entity which is defined by the input-output relation that the computer model executes, and can be seen as a function that transforms probability distributions over the set of possible inputs into probability distributions over the set of possible outputs.

It is shown here that both computer simulation and mathematical analysis are extremely useful tools to analyse this formal model, and they are certainly complementary in the sense that they can provide fundamentally different insights on the same model. Even more importantly, this chapter shows that there are plenty of synergies to be exploited by using the two techniques together.

The mathematical approach to analyse formal models consists in examining the rules that define the model directly. Its aim is to deduce the logical implications of these rules for any particular instance to which they can be applied. Our analysis of mathematical

---

L.R. Izquierdo (✉) • J.M. Galán • J.I. Santos  
Departamento de Ingeniería Civil, Universidad de Burgos, E-09001, Burgos, Spain  
e-mail: [luis@izquierdo.name](mailto:luis@izquierdo.name); [jmgalan@ubu.es](mailto:jmgalan@ubu.es); [jisantos@ubu.es](mailto:jisantos@ubu.es)

S.S. Izquierdo  
Departamento de Organización de Empresas y C.I.M., Universidad de Valladolid, E-47011,  
Valladolid, Spain  
e-mail: [segis@eis.uva.es](mailto:segis@eis.uva.es)

techniques to study formal models is focused on the theory of Markov Chains, which is particularly useful to characterise the dynamics of computer models.

In contrast with mathematical analysis, the computer simulation approach does not look at the rules that define the formal model directly, but instead tries to infer general properties of these rules by examining the outputs they produce when applied to particular instances of the input space. Thus, conclusions obtained with this approach may not be general. On a more positive note, computer simulation enables us to explore formal models beyond mathematical tractability, and we can achieve any arbitrary level of accuracy in our computational approximations by running the model sufficiently many times.

Bearing in mind the relative strengths and limitations of both approaches, this chapter explains three different ways in which mathematical analysis and computer simulation can be usefully combined to produce a better understanding of the dynamics of computer models. In doing so, it becomes clear that mathematical analysis and computer simulation should not be regarded as alternative – or even opposed – approaches to the formal study of social systems, but as complementary. Not only can they provide fundamentally different insights on the same model, but they can also produce hints for solutions for each other. In short, there are plenty of synergies to be exploited by using the two techniques together, so the full potential of each technique cannot be reached unless they are used in conjunction.

## 11.1 Introduction

This chapter is about how to better understand the dynamics of computer models using both simulation and mathematical analysis. Our starting point is a computer model which is already implemented and ready to be run; our objective is to gain a thorough understanding of its dynamics. Thus, this chapter is *not* about how to design, implement, verify, or validate a model; this chapter is about how to better understand its behaviour.

Naturally, we start by clearly defining our object of study: a computer model. The term ‘computer model’ can be understood in many different ways – i.e. seen from many different perspectives –, and not all of them are equally useful for every possible purpose. Thus, we start by interpreting the term ‘computer model’ in a way that will prove useful for our objective: to characterise and understand its behaviour. Once our object of study has been clearly defined, we then describe two techniques that are particularly useful to understand the dynamics of computer models: mathematical analysis and computer simulation.

In particular, this chapter will show that mathematical analysis and computer simulation should not be regarded as alternative – or even opposed – approaches to the formal study of social systems, but as complementary (Gotts et al. 2003a, b). They are both extremely useful tools to analyse formal models, and they are certainly complementary in the sense that they can provide fundamentally different insights on the same model. Even more importantly, this chapter will show that

there are plenty of synergies to be exploited by using the two techniques together, i.e. the full potential of each technique will not be reached until they are used in conjunction. The remaining of this introduction outlines the structure of the chapter.

Sections 11.2, 11.3 and 11.4 are devoted to explaining in detail what we understand by ‘computer model’, and they therefore provide the basic framework for the rest of the chapter. In particular, Sect. 11.2 shows that a computer model can be seen as an implementation – i.e. an explicit representation – of a certain deterministic input-output function in a particular programming language. This interpretation is very useful since, in particular, it will allow us to abstract from the details of the modelling platform where the computer model has been programmed, and focus on analysing the formal model that the computer model implements. This is clarified in Sect. 11.3, which explains that any computer model can be re-implemented in many different formalisms (in particular, in any sophisticated enough programming language), leading to alternative representations of the same input-output relation.

Most computer models in the Social Simulation literature make use of pseudo-random number generators. Section 11.4 explains that – for these cases and given our purposes – it is useful to abstract from the details of how pseudo-random numbers are generated, and look at the computer model as an implementation of a stochastic process. In a stochastic process, a certain input does not necessarily lead to one certain output only; instead, there are many different paths that the process may take with potentially different probabilities. Thus, in a stochastic process a certain input will generally lead to a particular probability distribution over the range of possible outputs, rather than to a single output only. Stochastic processes are used to formally describe how a system subjected to random events evolves through time.

Having explained our interpretation of the term ‘computer model’, Sect. 11.5 introduces and compares the two techniques to analyse formal models that are assessed in this chapter: computer simulation and mathematical analysis. The following two sections sketch possible ways in which each of these two techniques can be used to obtain useful insights about the dynamics of a model. Section 11.8 is then focused on the *joint use* of computer simulation and mathematical analysis. It is shown here that the two techniques can be used together to provide a picture of the dynamics of the model that could not be drawn by using one of the two techniques only. Finally, our conclusions are summarised in Sect. 11.9.

## 11.2 Computer Models as Input-Output Functions

At the most elementary level, a computer model can be seen as an implementation – i.e. an explicit representation – of a certain deterministic input-output function in a particular programming language. The word ‘function’ is useful because it

correctly conveys the point that any particular input given to the computer model will lead to one and only one output.<sup>1</sup> (Obviously, different inputs may lead to the same output.) Admittedly, however, the word ‘function’ may also mislead the reader into thinking that a computer model is necessarily simple. The computer model may be as complex and sophisticated as the programmer wants it to be but, ultimately, it is just an entity that associates a specific output to any given input, i.e. a function. In any case, to avoid confusion, we will use the term ‘formal model’ to denote the function that a certain computer model implements.<sup>2</sup> To be sure, the ‘formal model’ that a particular computer model implements is the abstract entity which is defined by the input-output relation that the computer model executes.<sup>3</sup>

Thus, running a computer model is just finding out the logical implications of applying a set of unambiguously defined formal rules (which are coded in the program and define the input-output function or formal model) to a set of inputs (Balzer et al. 2001). As an example, one could write the computer program “ $y = 4x$ ” and apply it to the input “ $x = 2$ ” to obtain the output “ $y = 8$ ”. The output ( $y = 8$ ), which is fully and unequivocally determined by the input ( $x = 2$ ) and the set of rules coded in the program ( $y = 4x$ ), can be seen as a theorem obtained by pure deduction ( $\{x = 2; y = 4x\} \Rightarrow y = 8$ ). Naturally, there is no reason why the inputs or the outputs should be numbers<sup>4</sup>; they could equally well be e.g. strings of characters. In the general case, a computer run is a logical theorem that reads: *the output obtained from running the computer simulation follows (with logical necessity) from applying to the input the algorithmic rules that define the model*. Thus, regardless of its inherent complexity, a computer run constitutes a perfectly valid sufficiency theorem (see e.g. Axtell 2000).

It is useful to realise that we could always apply the same inference rules ourselves to obtain – by logical deduction – the same output from the given input. While useful as a thought, when it comes to actually doing the job, it is much more convenient, efficient and less prone to errors to let computers derive the output for us. Computers are inference engines that are able to conduct many algorithmic processes at a speed that the human brain cannot achieve.

---

<sup>1</sup> Note that simulations of stochastic models are actually using pseudo-random number generators, which are deterministic algorithms that require a seed as an input.

<sup>2</sup> A formal model is a model expressed in a formal system (Cutland 1980). A formal system consists of a formal language and a deductive apparatus (a set of axioms and inference rules). Formal systems are used to derive new expressions by applying the inference rules to the axioms and/or previously derived expressions in the same system.

<sup>3</sup> The mere fact that the model has been implemented and can be run in a computer is a proof that the model is formal (Suber 2002).

<sup>4</sup> As a matter of fact, strictly speaking, inputs and outputs in a computer model are *never* numbers. We may interpret strings of bits as numbers, but we could equally well interpret the same strings of bits as e.g. letters. More importantly, a bit itself is already an abstraction, an interpretation we make of an electrical pulse that can be above or below a critical voltage threshold.

### 11.3 Different Ways of Representing the Same Formal Model

A somewhat controversial issue in the Social Simulation literature refers to the allegedly unique features of some modelling platforms. It is important to realise that any formal model implemented in a computer model can be re-implemented in many different programming languages, leading to exactly the same input-output relation. Different implementations are just different ways of representing one same formal model, much in the same way that one can say ‘Spain’ or ‘España’ to express the same concept in different languages: same thing, different representation, that’s all.

Thus, when analysing the dynamics of a computer model, it is useful to abstract from the details of the modelling platform that has been used to implement the computer model, and focus strictly on the formal model it represents, which could be re-implemented in any *sophisticated enough*<sup>5</sup> modelling platform. To be clear, let us emphasise that *any* computer model implemented in Objective-C (e.g. using Swarm) can be re-implemented in Java (e.g. using RePast or Mason), NetLogo, SDML, Mathematica© or Matlab©. Similarly, *any* computer model can be expressed as a well-defined mathematical function (Epstein 2006; Leombruni and Richiardi 2005; Richiardi et al. 2006).

Naturally, the implementation of a particular formal model may be more straightforward in some programming languages than in others. Programming languages differ in where they position themselves in the well-known trade-offs between ease of programming, functionality and performance; thus, different programming languages lead to more or less natural and more or less efficient implementations of any given formal model. Nonetheless, the important point is this: whilst we may have different implementations of the same formal model, and whilst each of these implementations may have different characteristics (in terms of e.g. code readability), ultimately they are all just different representations of the same formal model, and they will therefore return the same output when given the same input.

In the same way that using one or another formalism to represent a particular formal model will lead to more or less natural implementations, different formalisms also make more or less apparent certain properties of the formal model they implement. For example, we will see in this chapter that representing a computer model as a Markov chain, i.e. looking at the formal model implemented in a

---

<sup>5</sup> A sufficient condition for a programming language to be “sophisticated enough” is to allow for the implementation of the following three control structures:

- Sequence (i.e. executing one subprogram, and then another subprogram),
- Selection (i.e. executing one of two subprograms according to the value of a boolean variable, e.g. IF[boolean == true]-THEN[subprogram1]-ELSE[subprogram2]), and
- Iteration (i.e. executing a subprogram until a boolean variable becomes false, e.g. WHILE [boolean == true]-DO[subprogram]).

Any programming language that can combine subprograms in these three ways can implement any computable function; this statement is known as the “structured program theorem” (Böhm and Jacopini 1966; Harel 1980; Wikipedia 2007).

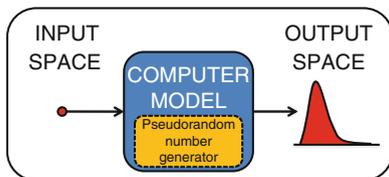
computer model through Markov's glasses, can make apparent various features of the computer model that may not be so evident without such glasses. In particular, as we will show later, Markov theory can be used to find out whether the initial conditions of a model determine its asymptotic dynamics or whether they are actually irrelevant in the long term. Also, the theory can reveal whether the model will sooner or later be trapped in an absorbing state.

## 11.4 'Stochastic' Computer Models as Stochastic Processes

Most computer models in the Social Simulation literature contain stochastic components. This section argues that, for these cases and given our purposes, it is convenient to revise our interpretation of computer models as *deterministic* input-output relations, abstract from the (deterministic) details of how pseudo-random numbers are generated, and reinterpret the term 'computer model' as an implementation of a stochastic process. This interpretation will prove useful in most cases and, importantly, does not imply any loss of generality: even if the computer model to be analysed does not contain any stochastic components, our interpretation will still be valid.

In the general case, the computer model to be analysed will make use of (what are meant to be) random numbers, i.e. the model will be stochastic. The word 'stochastic' requires some clarification. Strictly speaking, there does not exist a *truly stochastic* computer model, but one can approximate randomness to a very satisfactory extent by using pseudo-random number generators. The pseudo-random number generator is a deterministic algorithm that takes as input a value called the random seed, and generates a sequence of numbers that approximates the properties of random numbers. The sequence is not truly random in that it is completely determined by the value used to initialise the algorithm, i.e. the random seed. Therefore, if given the same random seed, the pseudo-random number generator will produce exactly the same sequence of (pseudo-random) numbers. (This fact is what made us define a computer model as an implementation of a certain *deterministic* input-output function in Sect. 11.2.)

Fortunately, the sequences of numbers provided by current off-the-shelf pseudo-random number generators approximate randomness remarkably well. This basically means that, for most intents and purposes in this discipline, it seems safe to assume that the pseudo-random numbers generated in one simulation run will follow the intended probability distributions to a satisfactory degree. The only problem we might encounter appears when running several simulations which we would like to be statistically independent. As mentioned above, if we used the same random seed for every run, we would obtain the same sequence of pseudo-random numbers, i.e. we would obtain exactly the same results. How can we *truly randomly* select a random seed? Fortunately, for most applications in this discipline, the state of the computer system at the time of starting a new run can be considered a truly random variable; and, conveniently, if no seed is explicitly provided to the pseudo-random number



**Fig. 11.1** A computer model can be usefully seen as the implementation of a function that transforms any given input into a certain probability distribution over the set of possible outputs

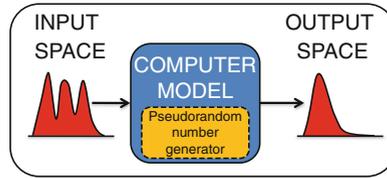
generator, most platforms generate a seed from the state of the computer system (e.g. using the time). When this is done, the sequences of numbers obtained with readily available pseudo-random number generators approximate statistical randomness and independence remarkably well.

Given that – for most intents and purposes in this discipline – we can safely assume that pseudo-random numbers are random and independent enough, we dispense with the qualifier ‘pseudo’ from now on for convenience. Since every random variable in the model follows a specific probability distribution, the computer model will indeed generate a particular probability distribution over the range of possible outputs. Thus, to summarise, a computer model can be usefully seen as the implementation of a stochastic process, i.e. a function that transforms any given input into a certain probability distribution over the set of possible outputs (Fig. 11.1).

Seeing that we can satisfactorily simulate random variables, note that studying the behaviour of a model that has been parameterised stochastically does not introduce any conceptual difficulties. In other words, we can study the behaviour of a model that has been parameterised with probability distributions rather than certain values. An example would be a model where agents start at a random initial location.

To conclude this section, let us emphasise an important corollary of the previous paragraphs: *any statistic that we extract from a parameterised computer model follows a specific probability distribution* (even if the values of the input parameters have been expressed as probability distributions).<sup>6</sup> Thus, a computer model can be seen as the implementation of a function that transforms probability distributions over the set of possible inputs into probability distributions over the set of possible outputs (Fig. 11.2). The rest of the chapter is devoted to characterising this function.

<sup>6</sup> Note that statistics extracted from the model can be of any nature, as long as they are unambiguously defined. For example, they can refer to various time-steps, and only to certain agents (e.g. “average wealth of female agents in odd time-steps from 1 to 99”).



**Fig. 11.2** A computer model can be seen as the implementation of a function that transforms probability distributions over the set of possible inputs into probability distributions over the set of possible outputs

## 11.5 Tools to Understand the Behaviour of Formal Models

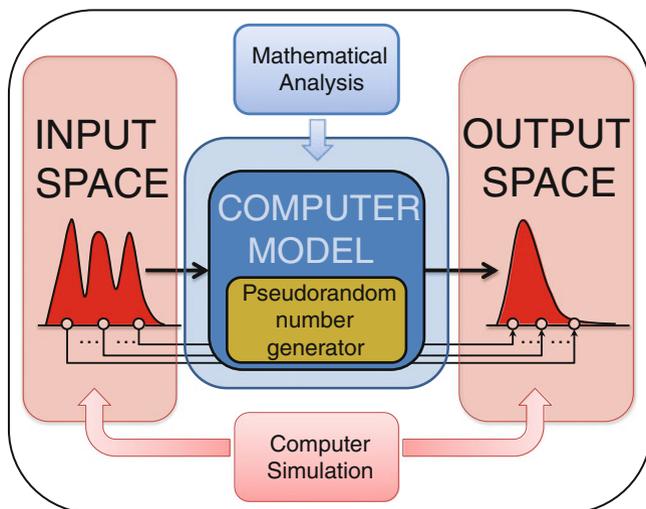
Once it is settled that a computer model can be seen as a particular implementation of a (potentially stochastic) function in a certain programming language, let us refer to such a function as the ‘formal model’ that the computer model implements. As mentioned before, this formal model can be expressed in many different formalisms – in particular, it can always be expressed as a set of well defined mathematical equations (Leombruni and Richiardi 2005) – and our objective consists in understanding its behaviour. To do that, we count with two very useful tools: mathematical analysis<sup>7</sup> and computer simulation.

The advantages and limitations of these two tools to formally study social systems have been discussed at length in the literature (see e.g. Axtell 2000; Axtell and Epstein 1994; Edmonds 2005; Gilbert 1999; Gilbert and Troitzsch 1999; Gotts et al. 2003a; Holland and Miller 1991; Ostrom 1988). Here we only highlight the most prominent differences between these two techniques (see Fig. 11.3).

In broad terms, when using mathematical analysis, one examines the rules that define the formal model directly, and tries to draw general conclusions about these rules. These conclusions are obtained by using logical deduction; hence they follow with logical necessity from the premises of the formal model (and the axioms of the mathematics employed). The aim when using mathematical analysis is usually to “solve” the formal system (or, most often, certain aspects of it) by producing general closed-form solutions that can be applied to *any* instance of the whole input set (or, at least, to large portions of the input set). Since the inferences obtained with mathematical analysis pertain to the rules themselves, such inferences can be safely particularised to any specific parameterisation of the model, even if such a parameterisation was never explicitly contemplated when analysing the model mathematically. This greatly facilitates conducting sensitivity analyses and assessing the robustness of the model.

Computer simulation is a rather different approach to the characterisation of the formal model (Epstein 2006; Axelrod 1997a). When using computer simulation,

<sup>7</sup> We use the term “mathematical analysis” in its broadest sense, i.e. we do not refer to any particular branch of mathematics, but to the general use of (any type of) mathematical technique to analyse a system.



**Fig. 11.3** In general terms, mathematical analysis tends to examine the rules that define the formal model directly. In contrast, computer simulation tries to infer general properties of such rules by looking at the outputs they produce when applied to particular instances of the input space

one often treats the formal model as a black box, i.e. a somewhat obscure abstract entity that returns certain outputs when provided with inputs. Thus, the path to understand the behaviour of the model consists in obtaining many input-output pairs and – using generalisation by induction – inferring general patterns about how the rules transform the inputs into the outputs (i.e. how the formal model works).

Importantly, the execution of a simulation run, i.e. the logical process that transforms any (potentially stochastic) given input into its corresponding (potentially stochastic) output is pure deduction (i.e. strict application of the formal rules that define the model). Thus, running the model in a computer provides a formal proof that a particular input (together with the set of rules that define the model) is sufficient to generate the output that is observed during the simulation. This first part of the computer simulation approach is therefore, in a way, very “mathematical”: outputs obtained follow with logical necessity from applying to the inputs the algorithmic rules that define the model.

In contrast, the second part of the computer simulation approach, i.e. inferring general patterns from particular instances of input-output pairs, can only lead to probable – rather than necessarily true – conclusions.<sup>8</sup> The following section explains how to rigorously assess the confidence we can place on the conclusions obtained using computer simulation, but the simple truth is irrefutable: inferences obtained using generalisation by induction can potentially fail when applied to instances that were not used to infer the general pattern. This is the domain of statistical extrapolation.

<sup>8</sup> Unless, of course, all possible particular instances are explored.

So why bother with computer simulation at all? The answer is clear: computer simulation enables us to study formal systems in ways that go beyond mathematical tractability. This role should not be underestimated: most models in the Social Simulation literature are mathematically intractable, and in such cases computer simulation is our only chance to move things forward. As a matter of fact, the formal models that many computer programs implement are often so complicated and cumbersome that the computer code itself is not that far from being one of the best descriptions of the formal model that can be provided.

Computer simulation can be very useful even when dealing with formal models that are mathematically tractable. Valuable uses of computer simulation in these cases include conducting insightful initial explorations of the model and presenting dynamic illustrations of its results.

And there is yet another important use of computer simulation. Note that understanding a formal model in depth requires identifying the parts of the model (i.e. the subset of rules) that are responsible for generating particular (sub)sets of results or properties of results. Investigating this in detail often involves changing certain subsets of rules in the model, so one can pinpoint which subsets of rules are necessary or sufficient to produce certain results. Importantly, changing subsets of rules can make the original model mathematically intractable and in such (common) cases, computer simulation is, again, our only hope. In this context, computer simulation can be very useful to produce counter-examples. This approach is very common in the literature of e.g. evolutionary game theory, where several authors (see e.g. Hauert and Doebeli 2004; Imhof et al. 2005; Izquierdo and Izquierdo 2006; Lieberman et al. 2009; Nowak and May 1992; Nowak and Sigmund 1992, 1993; Santos et al. 2006; Traulsen et al. 2006) resort to computer simulations to assess the implications of assumptions made in mathematically tractable models (e.g. the assumptions of “infinite populations” and “random encounters”).

It is important to note that the fundamental distinction between mathematical analysis and computer simulation as presented here is *not* about whether one uses pen and paper or computers to analyse formal models. We can follow either approach with or without computers, and it is increasingly popular to do mathematical analysis with computers. Recent advancements in symbolic computation have opened up a new world of possibilities to conduct mathematical analyses (using e.g. Mathematica©). In other words, nowadays it is perfectly possible to use computers to *directly* examine the rules that define a formal model (see Fig. 11.3).

Finally, as so often in life, things are not black or white, but involve some shade of grey. Similarly, most models are not tractable *or* intractable in mathematical terms; most often they are *partially* tractable. It is in these cases where an adequate combination of mathematical analysis and computer simulation is particularly useful. We illustrate this fact in Sect. 11.8, but first let us look at each technique separately. The following two sections provide some guidelines on how computer simulation (Sect. 11.6) and mathematical analysis (Sect. 11.7) can be usefully employed to analyse formal models.

## 11.6 Computer Simulation: Approximating the Exact Probability Distribution by Running the Model

The previous sections have argued that *any* statistic obtained from a (stochastically or deterministically) parameterised model follows a specific probability distribution. The statistic could be anything as long as it is unambiguously defined; in particular, it could refer to one or several time-steps, and to one or various subcomponents of the model. Ideally, one would like to calculate the exact probability distribution for the statistic using mathematical analysis, but this will not always be possible. In contrast, using computer simulation we will always be able to approximate this probability distribution to any arbitrary level of accuracy; this section provides basic guidelines on how to do that.

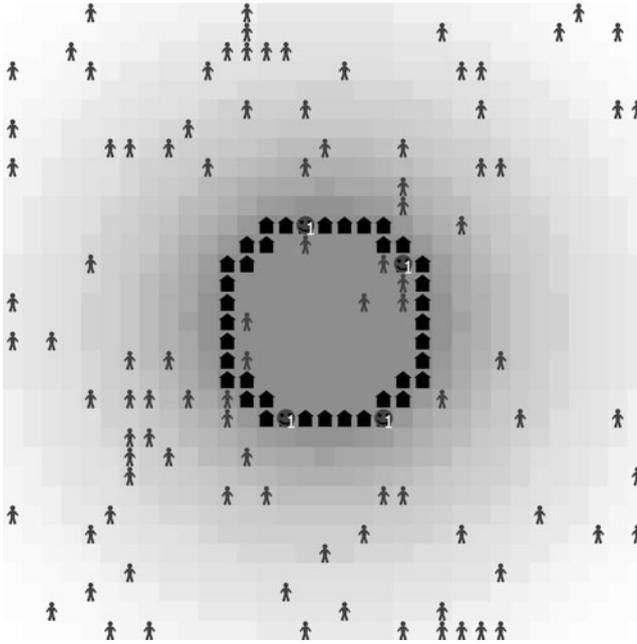
The output probability distribution – which is fully and unequivocally determined by the input distribution – can be approximated to any degree of accuracy by running enough simulation runs. Note that any specific simulation run will be conducted with a particular certain value for every parameter (e.g. a particular initial location for every agent), and will produce one and only one particular certain output (see Fig. 11.3). Thus, in order to infer the probability distribution over the set of outputs that a particular probability distribution over the set of inputs leads to, there will be a need to run the model many times (with different random seeds); this is the so-called Monte Carlo method.

The method is straightforward: obtain as many random samples as possible (i.e. run as many independent simulations as possible), since this will get us closer and closer to the exact distribution (by the law of large numbers). Having conducted a large number of simulation runs, the question that naturally comes to mind is: How close to the exact distribution is the one obtained by simulation?

To illustrate how to assess the quality of the approximation obtained by simulation, we use CoolWorld, a purpose-built agent-based model (Gilbert 2007) implemented in NetLogo 4.0 (Wilensky 1999). A full description of the model, an applet and the source code can be found at the dedicated model webpage <http://luis.izquierdo.name/models/coolworld>. For our purposes, it suffices to say that in CoolWorld there is a population of agents called walkers, who wander around a 2-dimensional grid made of square patches; some of the patches are empty whilst others contain a house (see Fig. 11.4). Patches are at a certain predefined temperature, and walkers tend to walk towards warmer patches, staying for a while at the houses they encounter in their journey.

Let us assume that we are interested in studying the number of CoolWorld walkers staying in a house in time-step 50. Initial conditions (which involve 100 walkers placed at a random location) are unambiguously defined at the model webpage and can be set in the implementation of CoolWorld provided by clicking on the button “Special conditions”. Figure 11.4 shows a snapshot of CoolWorld after having clicked on that button.

As argued before, given that the (stochastic) initial conditions are unambiguously defined, the number of CoolWorld walkers in a house after 50 time-steps will follow



**Fig. 11.4** Snapshot of CoolWorld. Patches are coloured according to their temperature: the higher the temperature, the darker the shade of *grey*. Houses are coloured in *black*, and form a circle around the central patch. Walkers are coloured in *grey*, and represented as a person if standing on a patch without a house, and as a *smiling face* if standing on a patch with a house. In the latter case, the *white label* indicates the number of walkers in the same house

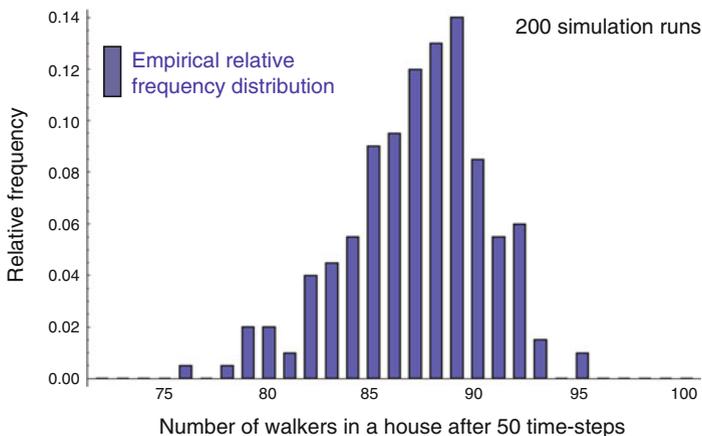
a specific probability distribution that we are aiming to approximate. For that, let us assume that we run 200 runs, and plot the relative frequency of the number of walkers in a patch with a house after 50 time-steps (see Fig. 11.5).

Figure 11.5 does not provide all the information that can be extracted from the data gathered. In particular, we can plot error bars showing the standard error for each calculated frequency without hardly any effort.<sup>9</sup> Standard errors give us information about the error we may be incurring when estimating the exact probabilities with the empirical frequencies. Another simple task that can be conducted consists in partitioning the set of runs into two batteries of approximately

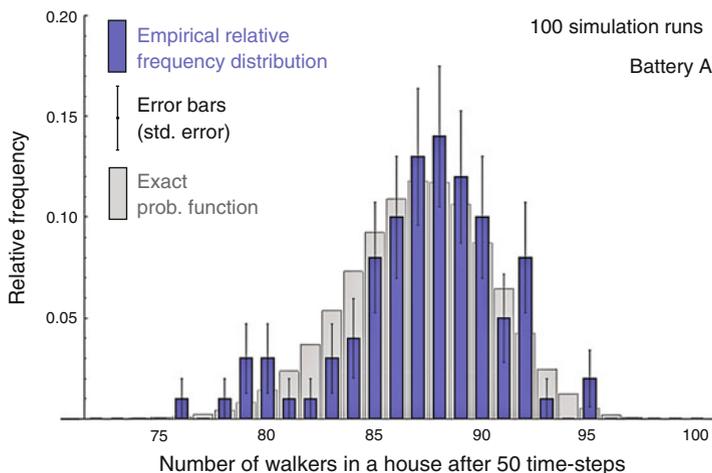
<sup>9</sup>The frequency of the event “there are  $i$  walkers in a patch with a house” calculated over  $n$  simulation runs can be seen as the mean of a sample of  $n$  i.i.d. Bernoulli random variables where success denotes that the event occurred and failure denotes that it did not. Thus, the frequency  $f$  is the maximum likelihood (unbiased) estimator of the exact probability with which the event occurs. The standard error of the calculated frequency  $f$  is the standard deviation of the sample divided by the square root of the sample size. In this particular case, the formula reads:

$$\text{Std. error } (f, n) = (f(1-f)/(n-1))^{1/2}$$

Where  $f$  is the frequency of the event,  $n$  is the number of samples, and the standard deviation of the sample has been calculated dividing by  $(n-1)$ .



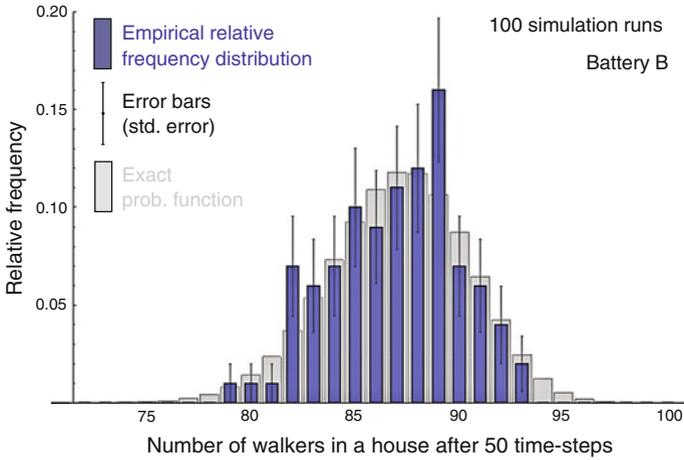
**Fig. 11.5** Relative frequency distribution of the number of walkers in a house after 50 time-steps, obtained by running CoolWorld 200 times, with initial conditions set by clicking on “Special conditions”



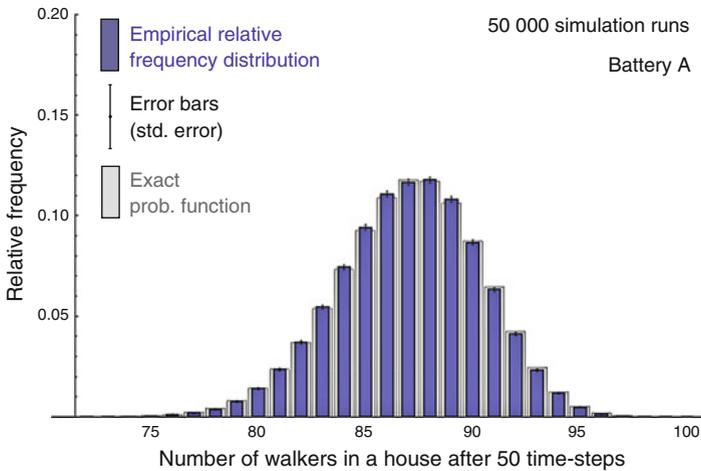
**Fig. 11.6** In *darker grey*: Relative frequency distribution of the number of walkers in a house after 50 time-steps, obtained by running CoolWorld 100 times (Battery A), with initial conditions set by clicking on “Special conditions”. In *lighter grey*: Exact probability distribution (Calculated using Markov chain analysis)

equal size and comparing the two distributions. If the two distributions are not similar, then there is no point in proceeding: we are not close to the exact distribution, so there is a need to run more simulations.

Figures 11.6 and 11.7 show the data displayed in Fig. 11.5 partitioned in two batteries of 100 simulation runs, including the standard errors. Figures 11.6 and 11.7 also show the exact probability distribution we are trying to approximate, which has been calculated using mathematical methods that are explained later in this chapter.

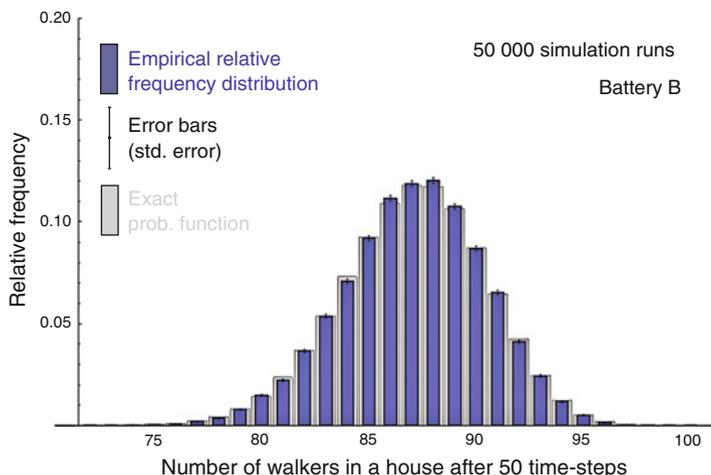


**Fig. 11.7** In *darker grey*: Relative frequency distribution of the number of walkers in a house after 50 time-steps, obtained by running CoolWorld 100 times (Battery B), with initial conditions set by clicking on “Special conditions”. In *lighter grey*: Exact probability distribution (Calculated using Markov chain analysis)



**Fig. 11.8** In *darker grey*: Relative frequency distribution of the number of walkers in a house after 50 time-steps, obtained by running CoolWorld 50,000 times (Battery A), with initial conditions set by clicking on “Special conditions”. In *lighter grey*: Exact probability distribution (Calculated using Markov chain analysis)

Figures 11.6 and 11.7 indicate that 100 simulation runs may not be enough to obtain a satisfactory approximation to the exact probability distribution. On the other hand, Figs. 11.8 and 11.9 show that running the model 50,000 times does seem to get us close to the exact probability distribution. The standard error, which



**Fig. 11.9** In *darker grey*: Relative frequency distribution of the number of walkers in a house after 50 time-steps, obtained by running CoolWorld 50,000 times (Battery B), with initial conditions set by clicking on “Special conditions”. In *lighter grey*: Exact probability distribution (Calculated using Markov chain analysis)

is inversely proportional to the square root of the sample size (i.e. the number of runs), is naturally much lower in these latter cases.

When, like in this example, the space of all possible outcomes in the distribution under analysis is finite (the number of walkers in a house must be an integer between 0 and 100), one can go further and calculate confidence intervals for the obtained frequencies. This is easily conducted when one realises that the exact probability distribution is a multinomial. Genz and Kwong (2000) show how to calculate these confidence intervals.

To conclude this section, let us emphasise that all that has been written here applies to *any* statistic obtained from *any* computer model. In particular, the statistic may refer to predefined regimes (e.g. “number of time-steps between 0 and 100 where there are more than 20 walkers in a house”) or to various time-steps (e.g. “total number of walkers in a house in odd time-steps in between time-steps 50 and 200”). These statistics, like any other one, follow a specific probability distribution that can be approximated to any degree of accuracy by running the computer model.

## 11.7 Mathematical Analysis: Time-Homogenous Markov Chains

The whole range of mathematical techniques that can be used to analyse formal systems is too broad to be reviewed here. Instead, we focus on one specific technique that seems to us particularly useful to analyse Social Simulation models:

Markov chain analysis. Besides, there are multiple synergies to be exploited by using Markov chain analysis and computer simulation together, as we will see in the next section.

Our first objective is to learn how to represent a particular computer model as a time-homogeneous Markov chain. This alternative representation of the model will allow us to use several simple mathematical results that will prove useful to understand the dynamics of the model. We therefore start by describing time-homogeneous Markov chains.

### 11.7.1 What Is a Time-Homogeneous Markov Chain?

Consider a system that in time-step  $n = \{1, 2, 3, \dots\}$  may be in one of a finite number of possible states  $S = \{s_1, s_2, \dots, s_M\}$ . The set  $S$  is called the state space; in this chapter we only consider *finite* state spaces.<sup>10</sup> Let the sequence of random variables  $X_n \in S$  represent the state of the system in time-step  $n$ . As an example,  $X_3 = s_9$  means that at time  $n = 3$  the system is in state  $s_9$ . The system starts at a certain initial state  $X_0$  and moves from one state to another. The system is stochastic in that, given the present state, the system may move to one or another state with a certain probability (see Fig. 11.10). The probability that the system moves from state  $i$  to state  $j$  in one time-step,  $P(X_{n+1} = j \mid X_n = i)$ , is denoted by  $p_{i,j}$ . As an example, in the Markov chain represented in Fig. 11.10,  $p_{4,6}$  equals 0 since the system cannot go from state 4 to state 6 in one single time-step. The system may also stay in the same state  $i$ , and this occurs with probability  $p_{i,i}$ . The probabilities  $p_{i,j}$  are called transition probabilities and they are often arranged in a matrix, namely the transition matrix  $P$ .

Implicitly, our definition of transition probabilities assumes two important properties about the system:

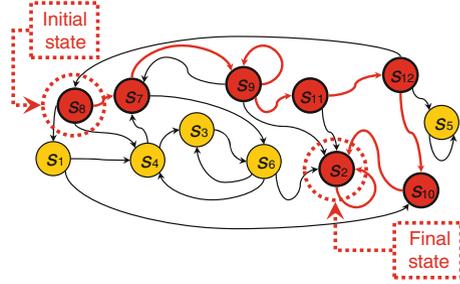
- (a) The system has the **Markov property**. This means that the present state contains all the information about the future evolution of the system that can be obtained from its past, i.e. given the present state of the system, knowing the past history about how the system reached the present state does not provide any additional information about the future evolution of the system. Formally,

$$P(X_{n+1} = x_{n+1} \mid X_n = x_n, X_{n-1} = x_{n-1}, \dots, X_0 = x_0) = P(X_{n+1} = x_{n+1} \mid X_n = x_n)$$

- (b) In this chapter we focus on **time-homogeneous** Markov chains, i.e. Markov chains with time-homogeneous transition probabilities. This basically means that transition probabilities  $p_{i,j}$  are independent of time, i.e. the one-step transition probability  $p_{i,j}$  depends on  $i$  and  $j$  but is the same at all times  $n$ . Formally,

<sup>10</sup>The term ‘Markov chain’ allows for countably infinite state spaces too (Karr 1990).

**Fig. 11.10** Schematic transition diagram of a Markov chain. *Circles* denote states and *directed arrows* indicate possible transitions between states. In this figure, thicker *circles* and *arrows* represent one possible path where the initial state  $X_0$  is  $s_8$  and the final state is  $s_2$



$$P(X_{n+1} = j | X_n = i) = P(X_n = j | X_{n-1} = i) = p_{i,j}$$

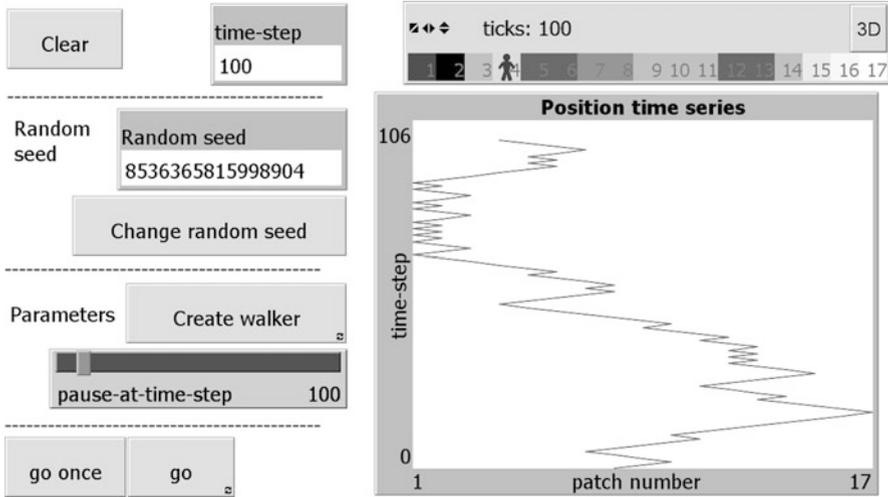
The crucial step in the process of representing a computer model as a time-homogeneous Markov chain (THMC) consists in identifying an appropriate set of state variables. A particular combination of specific values for these state variables will define one particular state of the system. Thus, the challenge consists in choosing the set of state variables in such a way that the computer model can be represented as a THMC. In other words, the set of state variables must be such that one can see the computer model as a transition matrix that unambiguously determines the probability of going from any state to any other state.

**Example: A Simple Random Walk**

Let us consider a model of a simple 1-dimensional random walk and try to see it as a THMC. In this model – which can be run and downloaded at the dedicated model webpage <http://luis.izquierdo.name/models/randomwalk> – there are 17 patches in line, labelled with the integers between 1 and 17. A random walker is initially placed on one of the patches. From then onwards, the random walker will move randomly to one of the spatially contiguous patches in every time-step (staying still is not an option). Space does not wrap around, i.e. patch 1’s only neighbour is patch 2 (Fig. 11.11).

This model can be easily represented as a THMC by choosing the agent’s position (e.g. the number of the patch she is standing on) as the only state variable. To be sure, note that defining the state of the system in this way, it is true that there is a fixed probability of going from any state to any other state, independent of time. The transition matrix  $P = [p_{i,j}]$  corresponding to the model is:

$$P = [p_{i,j}] = \begin{bmatrix} 0 & 1 & 0 & & & \dots & 0 \\ 0.5 & 0 & 0.5 & 0 & & & \vdots \\ 0 & 0.5 & 0 & 0.5 & 0 & & \\ & & 0 & 0.5 & 0 & 0.5 & 0 \\ & & & \ddots & \ddots & \ddots & \ddots \\ & & & & 0 & 0.5 & 0 & 0.5 & 0 \\ \vdots & & & & & & 0 & 0.5 & 0 & 0.5 \\ 0 & \dots & & & & & 0 & 1 & 0 \end{bmatrix} \tag{11.1}$$



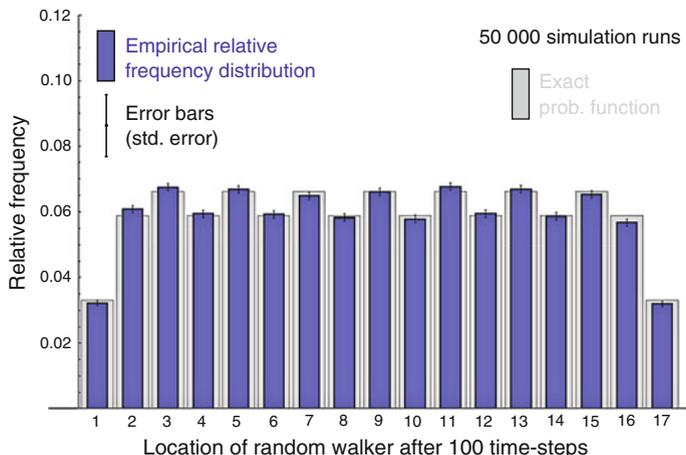
**Fig. 11.11** Snapshot of the 1-dimensional random walk applet. Patches are arranged in a horizontal line on the *top right corner* of the figure; they are labelled with integers, and coloured in shades of grey according to the number of times that the random walker has visited them: the higher the number of visits, the darker the shade of blue. The plot beneath the patches shows the time series of the random walker’s position

Where, as explained above,  $p_{i,j}$  is the probability  $P(X_{n+1} = j \mid X_n = i)$  that the system will be in state  $j$  in the following time-step, knowing that it is currently in state  $i$ .

### 11.7.2 Transient Distributions of Finite THMCs

The analysis of the dynamics of THMCs is usually divided into two parts: transient dynamics (finite time) and asymptotic dynamics (infinite time). The transient behaviour is characterised by the distribution of the state of the system  $X_n$  for a fixed time-step  $n \geq 0$ . The asymptotic behaviour (see Sects. 11.7.3 and 11.7.4) is characterised by the limit of the distribution of  $X_n$  as  $n$  goes to infinity, when this limit exists.

This section explains how to calculate the transient distribution of a certain THMC, i.e. the distribution of  $X_n$  for a fixed  $n \geq 0$ . In simple words, we are after a vector  $a^{(n)}$  containing the probability of finding the process in each possible state in time-step  $n$ . Formally,  $a^{(n)} = [a_1^{(n)}, \dots, a_M^{(n)}]$ , where  $a_i^{(n)} = P(X_n = i)$ , denotes the distribution of  $X_n$  for a THMC with  $M$  possible states. In particular,  $a^{(0)}$  denotes the initial distribution over the state space, i.e.  $a_i^{(0)} = P(X_0 = i)$ . Note that there is no problem in having uncertain initial conditions, i.e. probability functions over the space of possible inputs to the model.



**Fig. 11.12** Probability function of the position of the 1-dimensional random walker in time-step 100, starting at an initial random location

It can be shown that one can easily calculate the transient distribution in time-step  $n$ , simply by multiplying the initial conditions by the  $n$ -th power of the transition matrix  $P$ .

**Proposition 1.**  $a^{(n)} = a^{(0)} \cdot P^n$

Thus, the elements  $p^{(n)}_{i,j}$  of  $P^n$  represent the probability that the system is in state  $j$  after  $n$  time-steps having started in state  $i$ , i.e.  $p^{(n)}_{i,j} = P(X_n = j | X_0 = i)$ . A straightforward corollary of Proposition 1 is that  $a^{(n+m)} = a^{(n)} \cdot P^m$ .

As an example, let us consider the 1-dimensional random walk again. Imagine that the random walker starts at an initial random location, i.e.  $a^{(0)} = [1/17, \dots, 1/17]$ . The exact distribution of the walker’s position in time-step 100 would then be  $a^{(100)} = a^{(0)} \cdot P^{100}$ . This distribution is represented in Fig. 11.10, together with an empirical distribution obtained by running the model 50,000 times (Fig. 11.12).

Having obtained the probability function over the states of the system for any fixed  $n$ , namely the probability mass function of  $X_n$ , it is then straightforward to calculate the distribution of *any* statistic that can be extracted from the model. As argued in the previous sections, the state of the system fully characterises it, so *any* statistic that we obtain about the computer model in time-step  $n$  must be, ultimately, a function of  $\{X_0, X_1, \dots, X_n\}$ .

Admittedly, the transition matrix of most computer models cannot be easily derived, or it is unfeasible to operate with it. Nonetheless, this apparent drawback is not as important as one might expect. As we shall see below, it is often possible to infer many properties of a THMC even without knowing the exact values of its transition matrix, and these properties can yield useful insights about the dynamics of the associated process. Knowing the exact values of the transition matrix allows us to calculate the exact transient distributions using Proposition 1; this is desirable but not critical, since we can always approximate these distributions by conducting many simulation runs, as explained in Sect. 11.6.

### 11.7.3 Important Concepts

This section presents some basic concepts that will prove useful to analyse the dynamics of computer models. The notation used here follows the excellent book on stochastic processes written by Kulkarni (1995).

#### Definition 1: Accessibility

A state  $j$  is said to be accessible from state  $i$  if starting at state  $i$  there is a chance that the system may visit state  $j$  at some point in the future. By convention, every state is accessible from itself. Formally, a state  $j$  is said to be accessible from state  $i$  if for some  $n \geq 0$ ,  $p^{(n)}_{ij} > 0$ .

Note that  $j$  is accessible from  $i \neq j$  if and only if there is a directed path from  $i$  to  $j$  in the transition diagram. In that case, we write  $i \rightarrow j$ . If  $i \rightarrow j$  we also say that  $i$  leads to  $j$ . As an example, in the THMC represented in Fig. 11.10,  $s_2$  is accessible from  $s_{12}$  but not from  $s_5$ . Note that the definition of accessibility does not depend on the actual magnitude of  $p^{(n)}_{ij}$ , only on whether it is exactly zero or strictly positive.

#### Definition 2: Communication

A state  $i$  is said to communicate with state  $j$  if  $i \rightarrow j$  and  $j \rightarrow i$ .

If  $i$  communicates with  $j$  we also say that  $i$  and  $j$  communicate and write  $i \leftrightarrow j$ . As an example, note that in the simple random walk presented in Sect. 11.7.1, every state communicates with every other state. It is worth noting that the relation “communication” is transitive, i.e.

$$i \leftrightarrow j, j \leftrightarrow k \Rightarrow i \leftrightarrow k.$$

#### Definition 3: Communicating Class

A set of states  $C \subset S$  is said to be a communicating class if:

- Any two states in the communicating class communicate with each other. Formally,

$$i \in C, j \in C \Rightarrow i \leftrightarrow j$$

- The set  $C$  is maximal, i.e. no strict superset of a communicating class can be a communicating class. Formally,

$$i \in C, i \leftrightarrow j \Rightarrow j \in C$$

As an example, note that in the simple random walk presented in Sect. 11.7.1 there is one single communicating class that contains all the states. In the THMC represented in Fig. 11.10 there are four communicating classes:  $\{s_2\}$ ,  $\{s_5\}$ ,  $\{s_{10}\}$ ,  $\{s_1, s_3, s_4, s_6, s_7, s_8, s_9, s_{11}, s_{12}\}$ .

#### Definition 4: Closed Communicating Class (i.e. Absorbing Class). Absorbing State

A communicating class  $C$  is said to be closed if no state within  $C$  leads to any state outside  $C$ . Formally, a communicating class  $C$  is said to be closed if  $i \in C$  and  $j \notin C$  implies that  $j$  is not accessible from  $i$ .

Note that once a Markov chain visits a closed communicating class, it cannot leave it. Hence we will sometimes refer to closed communicating classes as “absorbing classes”. This latter term is not standard in the literature, but we find it useful here for explanatory purposes. Note that if a Markov chain has one single communicating class, it must be closed.

As an example, note that the communicating classes  $\{s_{10}\}$  and  $\{s_1, s_3, s_4, s_6, s_7, s_8, s_9, s_{11}, s_{12}\}$  in the THMC represented in Fig. 11.10 are not closed, as they can be abandoned. On the other hand, the communicating classes  $\{s_2\}$  and  $\{s_5\}$  are indeed closed, since they cannot be abandoned. When a closed communicating class consists of one single state, this state is called absorbing. Thus,  $s_2$  and  $s_5$  are absorbing states. Formally, state  $i$  is absorbing if and only if  $p_{i,i} = 1$  and  $p_{i,j} = 0$  for  $i \neq j$ .

**Proposition 2. Decomposition Theorem (Chung 1960)**

The state space  $S$  of any Markov chain can be uniquely partitioned as follows:

$$S = C_1 \cup C_2 \cup \dots \cup C_k \cup T$$

where  $C_1, C_2, \dots, C_k$  are closed communicating classes, and  $T$  is the union of all other communicating classes.

Note that we do not distinguish between non-closed communicating classes: we lump them all together into  $T$ . Thus, the unique partition of the THMC represented in Fig. 11.10 is  $S = \{s_2\} \cup \{s_5\} \cup \{s_1, s_3, s_4, s_6, s_7, s_8, s_9, s_{10}, s_{11}, s_{12}\}$ . The simple random walk model presented in Sect. 11.7.1 has one single (closed) communicating class  $C_1$  containing all the possible states, i.e.  $S \equiv C_1$ .

**Definition 5: Irreducibility**

A Markov chain is said to be irreducible if all its states belong to a single closed communicating class; otherwise it is called reducible. Thus, the simple random walk presented in Sect. 11.7.1 is irreducible, but the THMC represented in Fig. 11.10 is reducible.

**Definition 6: Transient and Recurrent States**

A state  $i$  is said to be transient if, given that we start in state  $i$ , there is a non-zero probability that we will never return back to  $i$ . Otherwise, the state is called recurrent. A Markov chain starting from a recurrent state will revisit it with probability 1, and hence revisit it infinitely often. On the other hand, a Markov chain starting from a transient state has a strictly positive probability of never coming back to it. Thus, a Markov chain will visit any transient state only finitely many times; eventually, transient states will not be revisited anymore.

**Definition 7: Periodic and Aperiodic States. Periodic and Aperiodic Communicating Classes**

A state  $i$  has period  $d$  if any return to state  $i$  must occur in multiples of  $d$  time-steps. If  $d = 1$ , then the state is said to be aperiodic; otherwise ( $d > 1$ ), the state is said to be periodic with period  $d$ . Formally, state  $i$ 's period  $d$  is the greatest common divisor of the set of integers  $n > 0$  such that  $p^{(n)}_{i,i} > 0$ . For our purposes, the

concept of periodicity is only relevant for recurrent states. As an example, note that every state in the simple random walk presented in Sect. 11.7.1 is periodic with period 2.

An interesting and useful fact is that if  $i \leftrightarrow j$ , then  $i$  and  $j$  must have the same period (see Theorem 5.2. in Kulkarni (1995)). In particular, note that if  $p_{i,i} > 0$  for any  $i$ , then the communicating class to which  $i$  belongs must be aperiodic. Thus, it makes sense to qualify communicating classes as periodic with period  $d$ , or aperiodic. A closed communicating class with period  $d$  can return to its starting state only at times  $d, 2d, 3d, \dots$ .

The concepts presented in this section will allow us to analyse the dynamics of any finite Markov chain. In particular, we will show that, given enough time, any finite Markov chain will necessarily end up in one of its closed communicating classes (i.e. absorbing classes).

### 11.7.4 Limiting Behaviour of Finite THMCs

This section is devoted to characterising the limiting behaviour of a THMC, i.e. studying the convergence (in distribution) of  $X_n$  as  $n$  tends to infinity. Specifically, we aim to study the behaviour of  $a_i^{(n)} = P(X_n = i)$  as  $n$  tends to infinity. From Proposition 1 it is clear that analysing the limiting behaviour of  $P^n$  would enable us to characterise  $a_i^{(n)}$ . There are many introductory books in stochastic processes that offer clear and simple methods to analyse the limiting behaviour of THMCs when the transition matrix  $P$  is tractable (see e.g. Chap. 5 in (Kulkarni 1999), Chaps. 2–4 in (Kulkarni 1995), Chap. 3 in (Janssen and Manca 2006) or the book chapter written by Karr (1990)). Nonetheless, we focus here on the general case, where operating with the transition matrix  $P$  may be computationally unfeasible.

#### 11.7.4.1 General Dynamics

The first step in the analysis of any THMC consists in identifying all the closed communicating classes, so we can partition the state space  $S$  as indicated by the decomposition theorem (see proposition 2). The following proposition (Theorems 3.7 and 3.8 in Kulkarni (1995)) reveals the significance of this partition:

#### **Proposition 3. General Dynamics of Finite THMCs**

Consider a finite THMC that has been partitioned as indicated in proposition 2. Then:

1. All states in  $T$  (i.e. not belonging to a closed communicating class) are transient.
2. All states in  $C_v$  (i.e. in any closed communicating class) are recurrent;  $v \in \{1, 2, \dots, k\}$ .

Proposition 3 states that sooner or later the THMC will enter one of the absorbing classes and stay in it forever. Formally, for all  $i \in S$  and all  $j \in T$ :

$\lim_{n \rightarrow \infty} p_{i,j}^{(n)} = 0$ , i.e. the probability of finding the process in a state belonging to a non-closed communicating class goes to zero as  $n$  goes to infinity. Naturally, if the initial state already belongs to an absorbing class  $C_v$ , then the chain will never abandon such a class. Formally, for all  $i \in C_v$  and all  $j \notin C_v$ :  $p_{i,j}^{(n)} = 0$  for all  $n \geq 0$ .

As an example of the usefulness of Proposition 3, consider the THMC represented in Fig. 11.10. This THMC has only two absorbing classes:  $\{s_2\}$  and  $\{s_5\}$ . Thus, the partition of the state space is:  $S = \{s_2\} \cup \{s_5\} \cup \{s_1, s_3, s_4, s_6, s_7, s_8, s_9, s_{10}, s_{11}, s_{12}\}$ . Hence, applying Proposition 3 we can state that the process will eventually end up in one of the two absorbing states,  $s_2$  or  $s_5$ . The probability of ending up in one or the other absorbing state depends on the initial conditions  $a^{(0)}$  (and on the actual numbers  $p_{i,j}$  in the transition matrix, of course). Slightly more formally, the limiting distribution of  $X_n$  exists, but it is not unique, i.e. it depends on the initial conditions.

#### 11.7.4.2 Dynamics Within Absorbing Classes

The previous section has explained that any simulation run will necessarily end up in a certain absorbing class; this section characterises the dynamics of a THMC that is already “trapped” in an absorbing class. This is precisely the analysis of irreducible Markov chains, since irreducible Markov chains are, by definition, Markov chains with one single closed communicating class (see definition 5). In other words, one can see any THMC as a set of transient states  $T$  plus a finite number of irreducible Markov sub-chains.

Irreducible THMCs behave significantly different depending on whether they are periodic or not. The following sections characterise these two cases.

#### Irreducible and Aperiodic THMCs

Irreducible and aperiodic THMCs are often called ergodic. In these processes the probability function of  $X_n$  approaches a limit as  $n$  tends to infinity. This limit is called the **limiting distribution**, and is denoted here by  $\pi$ . Formally, the following limit exists and is unique (i.e. independent of the initial conditions  $a_i^{(0)}$ ):

$$\lim_{n \rightarrow \infty} a_i^{(n)} = \pi_i > 0 \quad i \in S$$

Thus, in ergodic THMCs the probability of finding the system in each of its states in the long run is strictly positive and independent of the initial conditions (Theorems 3.7 and 3.15 in Kulkarni (1995)). As previously mentioned, calculating such probabilities may be unfeasible, but we can estimate them sampling many simulation runs at a sufficiently large time-step.

Importantly, in ergodic THMCs the limiting distribution  $\pi$  coincides with the **occupancy distribution**  $\pi^*$ , which is the long-run fraction of the time that the THMC spends in each state.<sup>11</sup> Naturally, the occupancy distribution  $\pi^*$  is also independent of the initial conditions. Thus, in ergodic THMCs, running just one simulation for long enough (which enables us to estimate  $\pi^*$ ) will serve to estimate  $\pi$  just as well.

The question that comes to mind then is: How long is long enough? I.e. when will I know that the empirical distribution obtained by simulation resembles the limiting distribution  $\pi$ ? Unfortunately there is no answer for that. The silver lining is that knowing that the limiting and the occupancy distribution coincide, that they must be stable in time, and that they are independent of the initial conditions, enables us to conduct a wide range of tests that may tell us when it is certainly *not* long enough. For example, we can run a battery of simulations and study the empirical distribution over the states of the system across samples as time goes by. If the distribution is not stable, then we have not run the model for long enough. Similarly, since the occupancy distribution is independent of the initial conditions, one can run several simulations with widely different initial conditions, and compare the obtained occupancy distributions. If the empirical occupancy distributions are not similar, then we have not run the model for long enough. Many more checks can be conducted.

Admittedly, when analysing a computer model one is often interested not so much in the distribution over the possible states of the system, but rather in the distribution of a certain statistic. The crucial point is to realise that if the statistic is a function of the state of the system (and all statistics that can be extracted from the model are), then the limiting and the occupancy distributions of the statistic exist, coincide and are independent of the initial conditions.

## Irreducible and Periodic THMCs

In contrast with aperiodic THMCs, the probability distribution of  $X_n$  in periodic THMCs does not approach a limit as  $n$  tends to infinity. Instead, in an irreducible THMC with period  $d$ , as  $n$  tends to infinity,  $X_n$  will in general cycle through  $d$  probability functions depending on the initial distribution.

As an example, consider the simple random walk again (which is irreducible and periodic, with period 2), and assume that the random walker starts at patch number 1 (i.e.  $X_0 = 1$ ). Given these settings, it can be shown that

---

<sup>11</sup> Formally, the occupancy of state  $i$  is defined as:

$$\pi_i^* = \lim_{n \rightarrow \infty} \frac{E(N_i(n))}{n + 1}$$

where  $N_i(n)$  denotes the number of times that the THMC visits state  $i$  over the time span  $\{0, 1, \dots, n\}$ .

$$\lim_{n \rightarrow \infty} a_i^{(2n)} = \left[ \frac{1}{16}, 0, \frac{1}{8}, 0, \frac{1}{8}, 0, \frac{1}{8}, 0, \frac{1}{8}, 0, \frac{1}{8}, 0, \frac{1}{8}, 0, \frac{1}{16} \right]$$

$$\lim_{n \rightarrow \infty} a_i^{(2n+1)} = \left[ 0, \frac{1}{8}, 0 \right]$$

In particular, the limits above show that the random walker cannot be at a patch with an even number in any even time-step, and he cannot be at a patch with an odd number in any odd time-step. In contrast, if the random walker started at patch number 2 (i.e.  $X_0 = 2$ ), then the limits above would be interchanged.

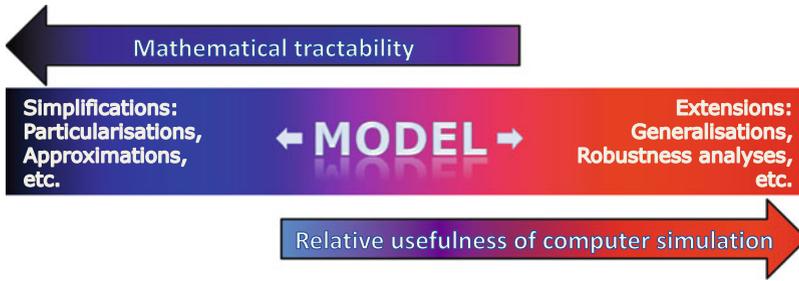
Fortunately, every irreducible (periodic or aperiodic) THMC does have a unique occupancy distribution  $\pi^*$ , independent of the initial conditions (see Theorem 5.19 in Kulkarni (1999)). In our particular example, this is:

$$\pi^* = \left[ \frac{1}{32}, \frac{1}{16}, \frac{1}{32} \right]$$

Thus, the long-run fraction of time that the system spends in each state in any irreducible THMC is unique (i.e. independent of the initial conditions). This is a very useful result, since any statistic which is a function of the state of the system will also have a unique occupancy distribution independent of the initial conditions. As explained before, this occupancy distribution can be approximated with one single simulation run, assuming it runs for long enough.

### 11.8 Synergies Between Mathematical Analysis and Computer Simulation

In this section we present various ways in which mathematical analysis and computer simulation can be combined to produce a better understanding of the dynamics of a model. Note that a full understanding of the dynamics of a model involves not only characterising (i.e. describing) them, but also finding out *why* such dynamics are being observed, i.e. identifying the subsets of rules that are necessary or sufficient to generate certain aspects of the observed dynamics. To do this, one often has to make changes in the model, i.e. build supporting models that differ only slightly from the original one and may yield useful insights about its dynamics. These supporting models will sometimes be more tractable (e.g. if heterogeneity or stochasticity are averaged out) and sometimes more complex (e.g. if interactions that were assumed to be global in the original model may only take place locally in the supporting model). Thus, for clarity, we distinguish three different cases and deal with them in turn (see Fig. 11.13):



**Fig. 11.13** To fully understand the dynamics of a model, one often has to study supporting models that differ only slightly from the original one. Some of these supporting models may be more tractable whilst others may be more complex

1. Characterisation of the dynamics of a model.
2. Moves towards greater mathematical tractability. This involves creating and studying supporting models that are simpler than the original one.
3. Moves towards greater mathematical complexity. This involves creating and studying supporting models that are less tractable than the original one.

### ***11.8.1 Characterising the Dynamics of the Model***

There are many types of mathematical techniques that can be usefully combined with computer simulation to characterise the dynamics of a model (e.g. Stochastic Approximation Theory (Benveniste et al. 1990; Kushner and Yin 1997)), but for limitations of space we focus here on Markov chain analysis only.

When using Markov chain analysis to characterise the dynamics of a model it may happen that the transition matrix can be easily computed and we can operate with it, or it may not. In the former case – which is quite rare in Social Simulation models –, one can provide a full characterisation of the dynamics of the model just by operating with the transition matrix (see Proposition 1 and the beginning of Sect. 11.7.4 for references). In general, however, deriving and operating with the transition matrix may be unfeasible, and it is in this common case where there is a lot to gain in using Markov chain analysis and computer simulation together. The overall method goes as follows:

- Use Markov chain analysis to assess the relevance of initial conditions and to identify the different regimes in which the dynamics of the model may end up trapped.
- Use the knowledge acquired in the previous point to design suitable computational experiments aimed at estimating the exact probability distributions for the relevant statistics (which potentially depend on the initial conditions).

The following describes this overall process in greater detail. Naturally, the first step consists in finding an appropriate definition of the state of the system, as explained in Sect. 11.7.1. The next step is to identify all the closed communicating (i.e. absorbing) classes in the model  $C_v$  ( $v \in \{1, 2, \dots, k\}$ ). This allows us to partition the state space of the Markov chain as the union of all the closed communicating classes  $C_1, C_2, \dots, C_k$  in the model plus another class  $T$  containing all the states that belong to non-closed communicating classes. Izquierdo et al. (2009) illustrate how to do this in ten well-known models in the Social Simulation literature.

In most cases, conducting the partition of the state space is not as difficult as it may seem at first. In particular, the following proposition provides some simple sufficient conditions that guarantee that the computer model contains one single aperiodic absorbing class, i.e. the finite THMC that the computer model implements is irreducible and aperiodic (i.e. ergodic).

**Proposition 4. Sufficient Conditions for Irreducibility and Aperiodicity**

1. If it is possible to go from any state to any other state in one single time-step ( $p_{i,j} > 0$  for all  $i \neq j$ ) and there are more than two states, then the THMC is irreducible and aperiodic.
2. If it is possible to go from any state to any other state in a finite number of time-steps ( $i \leftrightarrow j$  for all  $i \neq j$ ), and there is at least one state in which the system may stay for two consecutive time-steps ( $p_{i,i} > 0$  for some  $i$ ), then the THMC is irreducible and aperiodic.
3. If there exists a positive integer  $n$  such that  $p_{i,j}^{(n)} > 0$  for all  $i$  and  $j$ , then the THMC is irreducible and aperiodic (Janssen and Manca 2006, p. 107).

If one sees the transition diagram of the Markov chain as a (directed) network, the conditions above can be rewritten as:

1. The network contains more than two nodes and there is a directed link from every node to every other node.
2. The network is strongly connected and there is at least one loop.
3. There exists a positive integer  $n$  such that there is at least one walk of length  $n$  from any node to every node (including itself).

Izquierdo et al. (2009) show that many models in the Social Simulation literature satisfy one of these sufficient conditions (e.g. Epstein and Axtell's (1996) Sugarscape, Axelrod's (1986) metanorms models, Takahashi's (2000) model of generalized exchange, and Miller and Page's (2004) standing ovation model with noise). This is important since, as explained in Sect. 11.7.4.2, in ergodic THMCs the limiting and the occupancy distributions of any statistic exist, coincide and are independent of the initial conditions (so running just one simulation for long enough, which enables us to estimate the occupancy distribution, will serve to estimate the limiting distribution just as well).

Let us return to the general case. Having partitioned the state space, the analysis of the dynamics of the model is straightforward: all states in  $T$  (i.e. in any finite

communicating class that is not closed) are transient, whereas all states in  $C_v$  (i.e. in any finite closed communicating class) are recurrent. In other words, sooner or later any simulation run will enter one of the absorbing classes  $C_v$  and stay in it forever.

Here computer simulation can play a crucial role again, since it allows us to estimate the probability of ending up in each of the absorbing classes for any (stochastic or deterministic) initial condition we may be interested in. A case-in-point would be a model that has only a few absorbing states, or where various absorbing states are put together into only a few groups. Izquierdo et al. (2009) analyse models that follow that pattern: Axelrod's (1997b) model of dissemination of culture, Arthur's (1989) model of competing technologies, and Axelrod and Bennett's (1993) model of competing bimodal coalitions. CharityWorld (Polhill et al. 2006; Izquierdo and Polhill 2006) is an example of a model with a unique absorbing state.

The following step consists in characterising the dynamics of the system within each of the absorbing classes. Once the system has entered a certain absorbing class  $C_v$ , it will remain in it forever exhibiting a unique conditional<sup>12</sup> occupancy distribution  $\pi_{v,*}$  over the set of states that compose  $C_v$ . Naturally, the same applies to any statistic we may want to study, since all statistics that can be extracted from the model are a function of the state of the system.

The conditional occupancy distribution  $\pi_{v,*}$  denotes the (strictly positive) long-run fraction of the time that the system spends in each state of  $C_v$  given that the system has entered  $C_v$ . Importantly, the conditional occupancy distribution  $\pi_{v,*}$  is the same regardless of the specific state through which the system entered  $C_v$ . The role of simulation here is to estimate these conditional occupancy distributions for the relevant statistics by running the model for long enough.

Finally, recall that some absorbing classes are periodic and some are aperiodic. Aperiodic absorbing classes have a unique conditional limiting distribution  $\pi_v$ , denoting the long-run (strictly positive) probability of finding the system in each of the states that compose  $C_v$  given that the system has entered  $C_v$ . This conditional limiting distribution  $\pi_v$  coincides with the conditional occupancy distribution  $\pi_{v,*}$  and, naturally, is also independent of the specific state through which the system entered  $C_v$ . (Again, note that this also applies to the distribution of any statistic, as they are all functions of the state of the system, necessarily.)

In contrast with aperiodic absorbing classes, periodic absorbing classes do not generally have a unique limiting distribution; instead, they cycle through  $d$  probability functions depending on the specific state through which the system entered  $C_v$  (where  $d$  denotes the period of the periodic absorbing class). This is knowledge that one must take into account at the time of estimating the relevant probability distributions using computer simulation.

Thus, it is clear that Markov chain analysis and computer simulation greatly complement each other. Markov chain analysis provides the overall picture of the dynamics of the model by categorising its different dynamic regimes and

---

<sup>12</sup> Given that the system has entered the absorbing class  $C_v$ .

identifying when and how initial conditions are relevant. Computer simulation uses this information to design appropriate computational experiments that allow us to quantify the probability distributions of the statistics we are interested in. As explained above, these probability distributions can always be approximated with any degree of accuracy by running the computer model several times.

There are several examples of this type of synergetic combination of Markov chain analysis and computer simulation in the literature. Galán and Izquierdo (2005) analysed Axelrod's (1986) agent-based model as a Markov chain, concluded that the long-run behaviour of that model was independent of the initial conditions, in contrast to the initial conclusions of the original analysis. Galán and Izquierdo (2005) also used computer simulation to estimate various probability distributions. Ehrentreich (2002, 2006) used Markov chain analysis on the Artificial Stock Market (Arthur et al. 1997; LeBaron et al. 1999) to demonstrate that the mutation operator implemented in the model is not neutral to the learning rate, but introduces an upward bias.<sup>13</sup> A more positive example is provided by Izquierdo et al. (2007, 2008b), who used Markov chain analysis and computer simulation to confirm and advance various insights on reinforcement learning put forward by Macy and Flache (2002) and Flache and Macy (2002).

### ***11.8.2 Moves Towards Greater Mathematical Tractability: Simplifications***

There are at least two types of simplifications that can help us to better understand the dynamics of a model. One consists in studying specific parameterisations of the original model that are thought to lead to particularly simple dynamics, or to more tractable situations (Gilbert and Terna 2000; Gilbert 2007). Examples of this type of activity would be to run simulations without agents or with very few agents, explore the behaviour of the model using extreme parameter values, model very simple environments, etc. This activity is common practice in the field (see e.g. Gotts et al. 2003c, d).

A second type of simplification consists in creating an abstraction of the original model (i.e. a model of the model) which is mathematically tractable. An example of one possible abstraction would be to study the *expected* motion of the dynamic system (see the studies conducted by Galán and Izquierdo (2005), Edwards et al. (2003), Castellano et al. (2000), Huet et al. (2007), Mabrouk et al. (2007), Vilà (2008) and Izquierdo et al. (2007, 2008b) for illustrations of mean-field approximations). Since these mathematical abstractions do not correspond in a one-to-one way with the specifications of the formal model, any results obtained with them will not be conclusive in general, but they may give us insights

---

<sup>13</sup> This finding does not refute some of the most important conclusions obtained by the authors of the original model.

suggesting areas of stability and basins of attraction, clarifying assumptions, assessing sensitivity to parameters, or simply giving the option to illustrate graphically the expected dynamics of the original model. This approach can also be used as a verification technique to detect potential errors and artefacts (Galán et al. 2009).

### ***11.8.3 Moves Towards Greater Mathematical Complexity: Extensions***

As argued before, understanding the dynamics of a model implies identifying the set of assumptions that are responsible for particular aspects of the obtained results. Naturally, to assess the relevance of any assumption in a model, it is useful to replace it with other alternatives, and this often leads to greater mathematical complexity.<sup>14</sup>

Ideally, the evaluation of the significance of an assumption is conducted by generalisation, i.e. by building a more general model that allows for a wide range of alternative competing assumptions, and contains the original assumption as a particular case. An example would be the introduction of arbitrary social networks of interaction in a model where every agent necessarily interacts with every other agent. In this case, the general model with arbitrary networks of interaction would correspond with the original model if the network is assumed to be complete, but any other network could also be studied within the same common framework. Another example is the introduction of noise in deterministic models.

Building models by generalisation is useful because it allows for a transparent, structured and systematic way of exploring the impact of various alternative assumptions that perform the same role in the model, but it often implies a loss in mathematical tractability (see e.g. Izquierdo and Izquierdo 2006). Thus, it is often the case that a rigorous study of the impact of alternative assumptions in a model requires being prepared to slide up and down the tractability continuum depicted in Fig. 11.13 (Gotts et al. 2003b). In fact, all the cases that are mentioned in the rest of this section involved greater complexity than the original models they considered, and computer simulation had to be employed to understand their dynamics.

In the literature there are many examples of the type of activity explained in this section. For example, Klemm et al. studied the relevance of various assumptions in Axelrod's model of dissemination of culture (1997b) by changing the network topology (Klemm et al. 2003a), investigating the role of dimensionality (Klemm

---

<sup>14</sup> This is so because many assumptions we make in our models are, to some extent, for the sake of simplicity. As a matter of fact, in most cases the whole purpose of modelling is to build an abstraction of the world which is simpler than the world itself, so we can make inferences about the model that we cannot make directly from the real world (Edmonds 2001; Galán et al. 2009; Izquierdo et al. 2008a).

et al. 2003b, 2005), and introducing noise (Klemm et al. 2003c). Another example is given by Izquierdo and Izquierdo (2007), who analysed the impact of using different structures of social networks in the efficiency of a market with quality variability.

In the context of decision-making and learning, Flache and Hegselmann (1999) and Hegselmann and Flache (2000) compared two different decision-making algorithms that a set of players can use when confronting various types of social dilemmas. Similarly, Takadama et al. (2003) analysed the effect of three different learning algorithms within the same model.

Several authors, particularly in the literature of Game Theory, have investigated the effect of introducing noise in the decision-making of agents. This is useful not only to investigate the general effect of potential mistakes or experimentation, but also to identify the stochastic stability of different outcomes (see Sect. 10 in Izquierdo et al. (2009)). An illustrative example is given by Izquierdo et al. (2008b), who investigate the reinforcement learning algorithm proposed by Bush and Mosteller (1955) using both mathematical analysis and simulation, and find that the inclusion of small quantities of randomness in players' decisions can change the dynamics of the model dramatically.

Another assumption investigated in the literature is the effect of different spatial topologies (see e.g. Flache and Hegselmann (2001), who generalised two of their cellular automata models by changing their – originally regular – grid structure). Finally, as mentioned in Sect. 11.5, it is increasingly common in the field of evolutionary game theory to assess the impact of various assumptions using computer simulation (see e.g. Galán and Izquierdo 2005; Santos et al. 2006; Traulsen et al. 2006; Izquierdo and Izquierdo 2006).

## 11.9 Summary

In this chapter we have provided a set of guidelines to understand the dynamics of computer models using both simulation and mathematical analysis. In doing so, it has become clear that mathematical analysis and computer simulation should not be regarded as alternative – or even opposed – approaches to the formal study of social systems, but as complementary (Gotts et al. 2003a, b). Not only can they provide fundamentally different insights on the same model, but they can also produce hints for solutions for each other. In short, there are plenty of synergies to be exploited by using the two techniques together, so the full potential of each technique cannot be reached unless they are used in conjunction.

To understand the dynamics of any particular computer model, we have seen that it is useful to see the computer model as the implementation of a function that transforms probability distributions over the set of possible inputs into probability distributions over the set of possible outputs. We refer to this function as the formal model that the computer model implements.

The mathematical approach to analyse formal models consists in examining the rules that define the model directly; the aim is to deduce the logical implications of these rules for any particular instance to which they can be applied. Our analysis of mathematical techniques to study formal models has been focused on the theory of Markov Chains. This theory is particularly useful for our purposes since many computer models can be meaningfully represented as time-homogenous Markov chains.

In contrast with mathematical analysis, the computer simulation approach does not look at the rules that define the formal model directly, but instead tries to infer general properties of these rules by examining the outputs they produce when applied to particular instances of the input space. Thus, in the simulation approach, the data is produced by the computer using strict logical deduction, but the general patterns about how the rules transform the inputs into the outputs are inferred using generalisation by induction. Thus, in the general case – and in contrast with mathematical analysis –, the inferences obtained using computer simulation will not be necessarily correct in a strict logical sense; but, on the other hand, computer simulation enables us to explore formal models beyond mathematical tractability, and the confidence we can place on the conclusions obtained with this approach can be rigorously assessed in statistical terms. Furthermore, as shown in this chapter, we can achieve any arbitrary level of accuracy in our computational approximations by running the model sufficiently many times.

Bearing in mind the relative strengths and limitations of both approaches, we have identified at least three different ways in which mathematical analysis and computer simulation can be usefully combined to produce a better understanding of the dynamics of computer models.

The first synergy appears at the time of characterising the dynamics of the formal model under study. To do that, we have shown how Markov chain analysis can be used to provide an overall picture of the dynamics of the model by categorising its different dynamic regimes and identifying when and how initial conditions are relevant. Having conducted such an analysis, one can then use computer simulation to design appropriate computational experiments with the aim of quantifying the probability distributions of the variables we are interested in. These probability distributions can always be approximated with any degree of accuracy by running the computer model several times.

The two other ways in which mathematical analysis and computer simulation can be combined derive from the fact that understanding the dynamics of a model involves not only characterising (i.e. describing) them, but also finding out *why* such dynamics are being observed (i.e. discovering causality). This often implies building supporting models that can be simpler or more complex than the original one. The rationale to move towards simplicity is to achieve greater mathematical tractability, and this often involves studying particularly simple parameterisations of the original model, and creating abstractions which are amenable to mathematical analysis. The rationale to move towards complexity is to assess the relevance of specific assumptions, and it often involves building generalisations of the original model to explore the impact of competing assumptions that can perform the same role in the model but may lead to different results.

Let us conclude by encouraging the reader to put both mathematical analysis and computer simulation in their backpack, and be happy to glide up and down the tractability spectrum where both simple and complex models lie. The benefits are out there.

**Acknowledgements** The authors have benefited from the financial support of the Spanish Ministry of Education and Science (projects DPI2004-06590, DPI2005-05676 and TIN2008-06464-C03-02), the Spanish Ministry for Science and Innovation (CSD2010-00034), and the JCyL (projects VA006B09, BU034A08 and GREX251-2009). We are also very grateful to Nick Gotts, Bruce Edmonds, Gary Polhill, and Cesáreo Hernández for many extremely useful discussions.

## Further Reading

Firstly we suggest three things to read to learn more about Markov Chain models. Grinstead and Snell (1997) provides an excellent introduction to the theory of finite Markov Chains, with many examples and exercises. Häggström (2002) gives a clear and concise introduction to Probability theory and Markov Chain theory, and then illustrates the usefulness of these theories by studying a range of stochastic algorithms with important applications in optimisation and other problems in computing. One of the algorithms covered is the Markov chain Monte Carlo method. Finally, Kulkarni (1995) provides a rigorous analysis of many types of useful stochastic processes, e.g. discrete and continuous time Markov Chains, renewal processes, regenerative processes, and Markov regenerative processes.

The reader may find three other papers helpful. Izquierdo et al. (2009) analyses the dynamics of ten well-known models in the social simulation literature using the theory of Markov Chains, and is thus a good illustration of the approach in practice within the context of social simulation.<sup>15</sup> Epstein (2006) is a more general discussion, treating a variety of foundational and epistemological issues surrounding generative explanation in the social sciences, and discussing the role of agent-based computational models in generative social science. Finally, Leombruni and Richiardi (2005) usefully discusses several issues surrounding the interpretation of simulation dynamics and the generalisation of the simulation results. For a different approach to analysing the dynamics of a simulation model we refer the interested reader to Chap. 9 in this volume (Evans et al. 2013).

---

<sup>15</sup> This comment was added by the editors as the authors are too modest to so describe their own work.

## References

- Arthur WB (1989) Competing technologies, increasing returns, and lock-in by historical events. *Econ J* 99(394):116–131
- Arthur WB, Holland JH, LeBaron B, Palmer R, Tayler P (1997) Asset pricing under endogenous expectations in an artificial stock market. In: Arthur WB, Durlauf S, Lane D (eds) *The economy as an evolving complex system II*. Addison-Wesley Longman, Reading, pp 15–44
- Axelrod RM (1986) An evolutionary approach to norms. *Am Polit Sci Rev* 80(4):1095–1111
- Axelrod RM (1997a) Advancing the art of simulation in the social sciences. In: Conte R, Hegselmann R, Terna P (eds) *Simulating social phenomena (Lecture notes in economics and mathematical systems, 456)*. Springer, Berlin, pp 21–40
- Axelrod RM (1997b) The dissemination of culture: a model with local convergence and global polarization. *J Confl Resolut* 41(2):203–226
- Axelrod RM, Bennett DS (1993) A landscape theory of aggregation. *Br J Polit Sci* 23(2):211–233
- Axtell RL (2000) Why agents? On the varied motivations for agent computing in the social sciences. In: Macal VM, Sallach D (eds) *Proceedings of the workshop on agent simulation: applications, models, and tools*, Argonne National Laboratory, Argonne, pp 3–24
- Axtell RL, Epstein JM (1994) Agent based modeling: understanding our creations. *The Bulletin of the Santa Fe Institute*, Winter 1994, pp 28–32
- Balzer W, Brendel KR, Hofmann S (2001) Bad arguments in the comparison of game theory and simulation in social studies. *J Artif Soc Soc Simul* 4(2), <http://jasss.soc.surrey.ac.uk/4/2/1.html>
- Benveniste A, Métivier M, Priouret P (1990) *Adaptive algorithms and stochastic approximations*. Springer, Berlin
- Böhm C, Jacopini G (1966) Flow diagrams, turing machines and languages with only two formation rules. *Commun ACM* 9(5):366–371
- Bush RR, Mosteller F (1955) *Stochastic models for learning*. Wiley, New York
- Castellano C, Marsili M, Vespignani A (2000) Nonequilibrium phase transition in a model for social influence. *Phys Rev Lett* 85(16):3536–3539
- Cutland N (1980) *Computability: an introduction to recursive function theory*. Cambridge University Press, Cambridge
- Edmonds B (2001) The use of models: making MABS actually work. In: Moss S, Davidsson P (eds) *Multi-agent-based simulation (Lecture notes in artificial intelligence, 1979)*. Springer, Berlin, pp 15–32
- Edmonds B (2005) Simulation and complexity: how they can relate. In: Feldmann V, Mühlfeld K (eds) *Virtual worlds of precision: computer-based simulations in the sciences and social sciences*. Lit-Verlag, Münster, pp 5–32
- Edwards M, Huet S, Goreaud F, Deffuant G (2003) Comparing an individual-based model of behaviour diffusion with its mean field aggregate approximation. *J Artif Soc Soc Simul* 6(4), <http://jasss.soc.surrey.ac.uk/6/4/9.html>
- Ehrentreich N (2002) The Santa Fe Artificial Stock Market re-examined: suggested corrections (Betriebswirtschaftliche Diskussionsbeiträge Nr. 45/02). *Wirtschaftswissenschaftliche Fakultät, Martin-Luther-Universität, Halle/Saale*, <http://econwpa.wustl.edu:80/eps/comp/papers/0209/0209001.pdf>
- Ehrentreich N (2006) Technical trading in the Santa Fe Institute Artificial Stock Market revisited. *J Econ Behav Organ* 61(4):599–616
- Epstein JM (2006) Remarks on the foundations of agent-based generative social science. In: Judd KL, Tesfatsion L (eds) *Handbook of computational economics, vol. 2: agent-based computational economics*. North-Holland, Amsterdam, pp 1585–1604
- Epstein JM, Axtell RL (1996) *Growing artificial societies: social science from the bottom up*. Brookings Institution Press/MIT Press, Cambridge
- Evans A, Heppenstall A, Birkin M (2013) Understanding simulation results. Chapter 9 in this volume

- Flache A, Hegselmann R (1999) Rationality vs. learning in the evolution of solidarity networks: a theoretical comparison. *Comput Math Organ Theory* 5(2):97–127
- Flache A, Hegselmann R (2001) Do irregular grids make a difference? Relaxing the spatial regularity assumption in cellular models of social dynamics. *J Artif Soc Soc Simul* 4(4), <http://jasss.soc.surrey.ac.uk/4/4/6.html>
- Flache A, Macy MW (2002) Stochastic collusion and the power law of learning. *J Confl Resolut* 46(5):629–653
- Galán JM, Izquierdo LR (2005) Appearances can be deceiving: lessons learned re-implementing Axelrod's 'Evolutionary approach to norms'. *J Artif Soc Soc Simul* 8(3), <http://jasss.soc.surrey.ac.uk/8/3/2.html>
- Galán JM et al (2009) Errors and artefacts in agent-based modelling. *J Artif Soc Soc Simul* 12(1), <http://jasss.soc.surrey.ac.uk/12/1/1.html>
- Genz A, Kwong KS (2000) Numerical evaluation of singular multivariate normal distributions. *J Stat Comput Sim* 68(1):1–21
- Gilbert N (1999) Simulation: a new way of doing social science. *Am Behav Sci* 42(10):1485–1487
- Gilbert N (2007) Agent-based models, vol 153, Quantitative applications in the social sciences. Sage, London
- Gilbert N, Terna P (2000) How to build and use agent-based models in social science. *Mind Soc* 1(1):57–72
- Gilbert N, Troitzsch KG (1999) Simulation for the social scientist. Open University Press, Buckingham
- Gotts NM, Polhill JG, Law ANR (2003a) Agent-based simulation in the study of social dilemmas. *Artif Intell Rev* 19(1):3–92
- Gotts NM, Polhill JG, Adam WJ (2003b) Simulation and analysis in agent-based modelling of land use change. In: Online proceedings of the first conference of the European social simulation association, Groningen, The Netherlands, 18–21 Sept 2003, <http://www.uni-koblenz.de/~essa/ESSA2003/proceedings.htm>
- Gotts NM, Polhill JG, Law ANR (2003c) Aspiration levels in a land-use simulation. *Cybern Syst* 34(8):663–683
- Gotts NM, Polhill JG, Law ANR, Izquierdo LR (2003d) Dynamics of imitation in a land use simulation. In: Dautenhahn K, Nehaniv C (eds) Proceedings of the second international symposium on imitation in animals and artefacts, University of Wales, Aberystwyth, 7–11 April 2003, pp 39–46
- Grinstead CM, Snell JL (1997) Chapter 11: Markov chains. In: Grinstead CM, Snell JL (eds) Introduction to probability (Second revised edition). American Mathematical Society, Providence, pp 405–470, [http://www.dartmouth.edu/~chance/teaching\\_aids/books\\_articles/probability\\_book/book.html](http://www.dartmouth.edu/~chance/teaching_aids/books_articles/probability_book/book.html)
- Hägström O (2002) Finite Markov chains and algorithmic applications. Cambridge University Press, Cambridge
- Harel D (1980) On folk theorems. *Commun ACM* 23(7):379–389
- Hauer C, Doebeli M (2004) Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature* 428(6983):643–646
- Hegselmann R, Flache A (2000) Rational and adaptive playing. *Anal Krit* 22(1):75–97
- Holland JH, Miller JH (1991) Artificial adaptive agents in economic theory. *Am Econ Rev* 81(2): 365–370
- Huet S, Edwards M, Deffuant G (2007) Taking into account the variations of neighbourhood sizes in the mean-field approximation of the threshold model on a random network. *J Artif Soc Soc Simul* 10(1), <http://jasss.soc.surrey.ac.uk/10/1/10.html>
- Imhof LA, Fudenberg D, Nowak MA (2005) Evolutionary cycles of cooperation and defection. *Proc Natl Acad Sci USA* 102(31):10797–10800
- Izquierdo SS, Izquierdo LR (2006) On the structural robustness of evolutionary models of cooperation. In: Corchado E, Yin H, Botti VJ, Fyfe C (eds) Intelligent data engineering and automated learning – IDEAL 2006 (Lecture notes in computer science, 4224). Springer, Berlin/Heidelberg, pp 172–182

- Izquierdo SS, Izquierdo LR (2007) The impact on market efficiency of quality uncertainty without asymmetric information. *J Bus Res* 60(8):858–867
- Izquierdo LR, Polhill JG (2006) Is your model susceptible to floating point errors? *J Artif Soc Soc Simul* 9(4), <http://jasss.soc.surrey.ac.uk/9/4/4.html>
- Izquierdo LR, Izquierdo SS, Gotts NM, Polhill JG (2007) Transient and asymptotic dynamics of reinforcement learning in games. *Game Econ Behav* 61(2):259–276
- Izquierdo LR, Galán JM, Santos JI, Olmo R (2008a) Modelado de sistemas complejos mediante simulación basada en agentes y mediante dinámica de sistemas. *Empiria* 16:85–112
- Izquierdo SS, Izquierdo LR, Gotts NM (2008b) Reinforcement learning dynamics in social dilemmas. *J Artif Soc Soc Simul* 11(2), <http://jasss.soc.surrey.ac.uk/11/2/1.html>
- Izquierdo LR, Izquierdo SS, Galán JM, Santos JI (2009) Techniques to understand computer simulations: Markov chain analysis. *J Artif Soc Soc Simul* 12(1), <http://jasss.soc.surrey.ac.uk/12/1/6.html>
- Janssen J, Manca R (2006) *Applied semi-Markov processes*. Springer, New York
- Karr AF (1990) Markov processes. In: Heyman DP, Sobel MJ (eds) *Stochastic models*, vol 2, *Handbooks in operations research and management science*. Elsevier, Amsterdam, pp 95–123
- Klemm K, Eguíluz VM, Toral R, San Miguel M (2003a) Nonequilibrium transitions in complex networks: a model of social interaction. *Phys Rev E* 67(2):026120
- Klemm K, Eguíluz VM, Toral R, San Miguel M (2003b) Role of dimensionality in Axelrod's model for the dissemination of culture. *Phys A* 327(1–2):1–5
- Klemm K, Eguíluz VM, Toral R, San Miguel M (2003c) Global culture: a noise-induced transition in finite systems. *Phys Rev E* 67(4):045101
- Klemm K, Eguíluz VM, Toral R, San Miguel M (2005) Globalization, polarization and cultural drift. *J Econ Dyn Control* 29(1–2):321–334
- Kulkarni VG (1995) *Modeling and analysis of stochastic systems*. Chapman & Hall/CRC, Boca Raton
- Kulkarni VG (1999) *Modeling, analysis, design, and control of stochastic systems*. Springer, New York
- Kushner HJ, Yin GG (1997) *Stochastic approximation algorithms and applications*. Springer, New York
- Lebaron B, Arthur WB, Palmer R (1999) Time series properties of an artificial stock market. *J Econ Dyn Control* 23(9–10):1487–1516
- Leombruni R, Richiardi M (2005) Why are economists sceptical about agent-based simulations? *Phys A* 355(1):103–109
- Lieberman E, Havlin S, Nowak MA (2009) Evolutionary dynamics on graphs. *Nature* 433(7023):312–316
- Mabrouk N, Deffuant G, Lobry C (2007) Confronting macro, meso and micro scale modelling of bacteria dynamics. In: *M2M 2007: Third international model-to-model workshop*, Marseille, France, 15–16 Mar 2007, <http://m2m2007.macaulay.ac.uk/M2M2007-Mabrouk.pdf>
- Macy MW, Flache A (2002) Learning dynamics in social dilemmas. *Proc Natl Acad Sci USA* 99(3):7229–7236
- Miller JH, Page SE (2004) The standing ovation problem. *Complexity* 9(5):8–16
- Nowak MA, May RM (1992) Evolutionary games and spatial chaos. *Nature* 359(6398):826–829
- Nowak MA, Sigmund K (1992) Tit for tat in heterogeneous populations. *Nature* 355(6357):250–253
- Nowak MA, Sigmund K (1993) A strategy of win-stay, lose-shift that outperforms tit for tat in the Prisoner's Dilemma game. *Nature* 364(6432):56–58
- Ostrom T (1988) Computer simulation: the third symbol system. *J Exp Soc Psychol* 24(5):381–392
- Polhill JG, Izquierdo LR, Gotts NM (2006) What every agent based modeller should know about floating point arithmetic. *Environ Model Software* 21(3):283–309
- Richiardi M, Leombruni R, Saam NJ, Sonnessa M (2006) A common protocol for agent-based social simulation. *J Artif Soc Soc Simul* 9(1), <http://jasss.soc.surrey.ac.uk/9/1/15.html>

- Santos FC, Pacheco JM, Lenaerts T (2006) Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proc Natl Acad Sci USA* 103(9):3490–3494
- Suber P (2002) Formal systems and machines: an isomorphism (Electronic hand-out for the course “Logical Systems”, Earlham College, Richmond, IN) <http://www.earlham.edu/~peters/courses/logsys/machines.htm>
- Takadama K, Suematsu YL, Sugimoto N, Nawa NE, Shimohara K (2003) Cross-element validation in multiagent-based simulation: switching learning mechanisms in agents. *J Artif Soc Soc Simul* 6(4), <http://jasss.soc.surrey.ac.uk/6/4/6.html>
- Takahashi N (2000) The emergence of generalized exchange. *Am J Sociol* 10(4):1105–1134
- Traulsen A, Nowak MA, Pacheco JM (2006) Stochastic dynamics of invasion and fixation. *Phys Rev E* 74(1):011909
- Vilà X (2008) A model-to-model analysis of Bertrand competition. *J Artif Soc Soc Simul* 11(2), <http://jasss.soc.surrey.ac.uk/11/2/11.html>
- Wikipedia (2007) Structured program theorem, [http://en.wikipedia.org/w/index.php?title=Structured\\_program\\_theorem&oldid=112885072](http://en.wikipedia.org/w/index.php?title=Structured_program_theorem&oldid=112885072)
- Wilensky U (1999) NetLogo, <http://ccl.northwestern.edu/netlogo>