Mimicking Human Problem-Solving with Agents: Exploring Model Calibration

Dr. Anthony H. Dekker
Joint Operations Division
Defence Science and Technology Organisation (DSTO)
Canberra, Australia
dekker@ACM.org

Abstract. Agent-based social simulation models have a wide range of applications, and can incorporate numeric parameters of various kinds. In this paper, we use a simple agent-based simulation model of a laboratory experiment in network colouring to explore the selection of such numeric parameters. In particular, we examine two fundamental approaches to selecting model parameters (model calibration) based on empirical data: directly, comparing the data to model parameters; and indirectly, by comparing the data to model outputs. Using our model, we examine the strengths and weaknesses of the second approach. We discuss the insights provided by the model, and the extent to which confidence in these insights is justified when parameters are selected indirectly. The indirect approach to parameter selection has value in building social agent-based models, particularly when data on parameter values is unavailable, provided the number of parameters is relatively small.

1. INTRODUCTION

Agent-based social simulation models have practical applications in training and concept exploration in areas such as defence, emergency management, transportation, and economics. In this paper, we use a case study to examine the selection of numeric parameters in social agent-based models – a process called **model calibration** (Epstein 2008), which is essential if models are to give meaningful results, but which also poses challenges, especially for numeric parameters.

Agent-based social simulation models can incorporate several kinds of numeric parameter. Some parameters are simply **integer counts**. For example, a simulation of a school might be parameterised with the number of children per classroom. In that case, appropriate parameter values can usually be selected based on surveys, observations, or other fieldwork. Social agent-based models can also contain parameters which are **probabilities**. For example, models may include a probability that two nearby agents will interact in a certain way. Obtaining realistic probabilities is more difficult than obtaining counts, but nevertheless observations and other fieldwork can often still be used, as in Dekker (2009).

Parameters representing **physical variables**, such as the maximum speed of a vehicle, can often simply be measured in the real world, as in Crooks *et al*. (2009).

The most difficult type of numeric parameter is one representing a **psychological** or **social variable**. For example, an agent-based model may include a parameter representing the average level of trust between two people. In this case, we can obtain numeric values from the population under study – for example, by administering a survey where people rate trust on a Likert scale (Cohen *et al*. 1988) – but the process of translating between such empirical values and values within the model is far from obvious. The relationship between numeric model parameters and measured real-world parameters will often be nonlinear, and finding the correspondence between the two may be difficult.

For many applications, these difficulties are not actually of great concern. There are many reasons for modelling social systems, including guiding data collection, illuminating dynamics, and demonstrating trade-offs (Epstein 2008). For several of these goals, it is not necessary to select numeric parameters which accurately represent a particular real-world situation. Instead, it is often sufficient to sweep through the range of values – for example, from total trust to complete distrust – in order to explore the range of different behaviour patterns which result. Examples of such parameter explorations are given by Dekker (2007) and Özman (2008).

For applications where we wish to model a particular real-world situation, there are two fundamental approaches to selecting model parameters, as shown in Figure 1. In the first approach, parameter values are selected on the basis of empirical data (Gilbert 2008). This is particularly appropriate for parameters which are probabilities or integer counts. When this approach is used, comparison of the overall model behaviour with the real world can be used to validate both the overall simulation design and the choice of parameters. However, the empirical data needed to apply this approach is often unavailable.

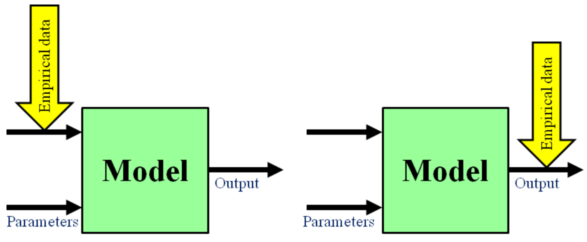

(a) Compare data to parameters (b) Compare data to output

Figure 1: Two ways of using empirical data to choose model parameters

The second approach to selecting numeric parameters is an indirect one (Moss 2008, paragraphs 2.8 and 2.9). In this case, empirical data is available for comparison with the output of the model. The full range of parameter values is explored, and values are selected which produce behaviour best matching the empirical data. In some cases, Bayesian methods can be used to make the selection (Windrum *et al*. 2007, paragraph 4.10). In the remainder of this paper, we describe an example of this approach, relating to a problem of group coordination, where a group of people solve a problem based on network colouring (Gibbons 1985).

For both approaches to parameter selection, and especially for the second, the task of selecting parameters becomes more difficult as the number of parameters increases, so that models with an excessive number of parameters are to be avoided.

This second approach requires some care. In particular, since the model outputs are used to select parameter values, comparison of the model outputs to the real world cannot also be used as a validation technique. This modelling approach describes a family of parameterised models $M(x)$, $M(x,y)$, or $M(x,y,z)$, etc. and allows us to choose the member of the family which best fits the real-world data. It does not, of itself, allow us to conclude that the internals of the model correspond to the real world, or even that the model parameters correspond to empirically observable ones (Windrum *et al*. 2007, paragraph 4.8). On the other hand, if there is independent evidence that the selected family of models is constructed in a way which corresponds to real-world processes, then the designers can have greater confidence that the selected parameter values correspond to variables in the real world.

This paper uses a simple case study to examine these issues, and in discussing the case study, we explore further the question of when such confidence is justified.

2. NETWORK COLOURING

Our case study is based on the network colouring problem (Gibbons 1985). The challenge here is to take a network, like the ones in Figure 2, and ensure that linked nodes have different colours, using the smallest possible set of colours to do so. For the networks in Figure 2, two colours are always sufficient. Our focus is on networks where each node is a person, who has to select their own colour, so that the problem is one of group coordination. Each person has to eliminate conflicts, which occur when linked people have the same colour. Furthermore, this is a problem with limited information, since people can only see the colours of those linked to them. As a human group activity, this problem was formulated and explored by Kearns *et al*. (2006), who collected empirical data on the time taken to solve the colouring problem by groups of people connected in different ways.

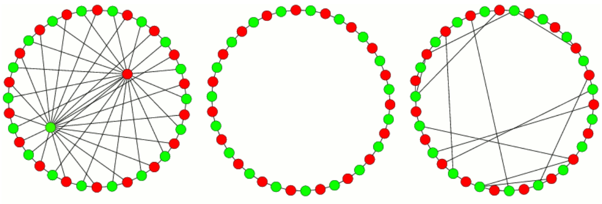
(a) Leader Cycle (b) Simple Ring (c) Twelve-Chord Ring

Figure 2: Three sample networks used in our colouring experiments, based on Kearns *et al*. (2006).

Kearns *et al*. (2006) used a graphical user interface tool, similar to the one shown in Figure 3, to allow participants in their experiments to visualise the colours selected by other linked participants. Furthermore, the tool provided buttons with which participants could select a new colour.

We have implemented an agent-based model for this problem. Figure 3 and Figure 4 show two different versions of the implementation, one of which allows interaction with a single human participant, while the other involves only software agents. Our goal has been to develop agents which tackle the network colouring problem in the same way as human beings, not necessarily in the most efficient possible way. We do this by using the empirical data collected by Kearns *et al*. (2006) to select parameters for the model. This case study is part of a larger activity to develop realistic agent-based models which can interact with human beings in an experimental or game setting (Dekker 2008), and we intend to use the results of this case study to inform that activity.

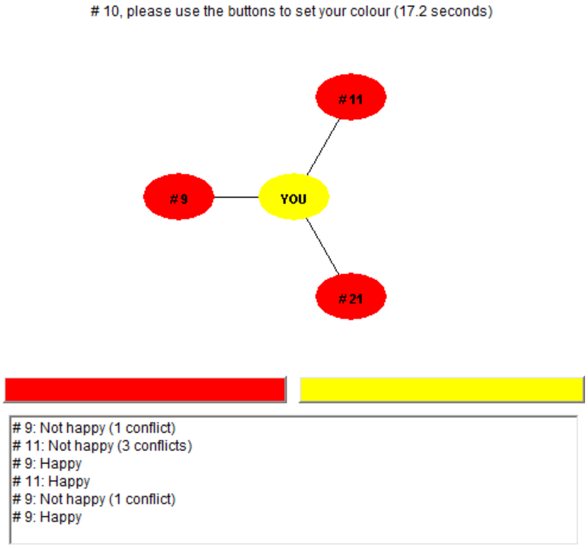
10, please use the buttons to set your colour (17.2 seconds)

9: Not happy (1 conflict)
11: Not happy (3 conflicts)
9: Happy
11: Happy
9: Not happy (1 conflict)
9: Happy

Figure 3: A graphical user interface tool for the network colouring problem, similar to the one used by Kearns *et al*. (2006). However, this version of the tool allows a single human participant to interact with a network of software agents, developed according to the method outlined here. At this point in time, the human participant has no conflicts with the three linked agents. The text communication by agents (shown below the two colour-selection buttons) was not part of the original study of Kearns *et al*.

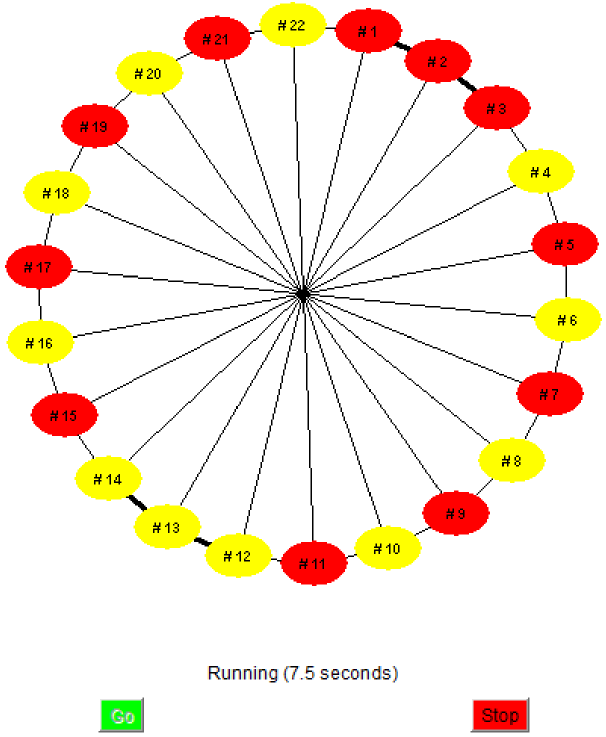

Figure 4: An agent implementation of the network colouring problem in which agents collaborate with each other, with no human participation. At the indicated point in time, there are four conflicts, indicated by thicker lines on the periphery of the network. These conflicts would be resolved if agents #2 and #13 changed their colours, and other agents retained theirs.

3. EXPERIMENTAL RESULTS

Kearns *et al.* (2006) note that the time taken for a group of people to resolve colouring conflicts depends on the number of colours required, and so we focus our attention on networks where only two colours are needed. Figure 2 illustrates the networks we used. Our twelve-chord ring represents an intermediate network between the two cross-linked rings (five-chord ring and twenty-chord ring) used by Kearns *et al.* (2006). Table 1 summarises the timing data collected by Kearns *et al.* The time taken by groups of people to solve the problem (in seconds) was roughly 20 times the average distance between nodes in the network. This reflects the fact that solving the problem required information to travel through the network, as individuals reacted to changes made by other participants. This process takes more time in networks with a greater average distance. From the data in Table 1, we predict that a group of people connected according to the twelve-chord ring in Figure 2 would take around 82.6 seconds to solve the problem. Expressing the data in Table 1 another way, we expect a group of people connected in a simple ring to take about 4.23 times as long to solve the problem as a group of people connected in a "leader cycle." Similarly, we expect a group of people connected in a twelve-chord ring to take 1.79 times as long as a group of people connected in a "leader cycle."

Table 1: Empirical data for the four ring-like networks used for human experiments by Kearns *et al.* (2006) and for our three networks from Figure 2.

Network	Average distance D	Human time taken in seconds (T)	Ratio T/D	Predicted time (20 D)
Simple Ring	9.76	144.2 (underestimate, only 5 groups of 6 completed the task)	14.8	195.2
Five-Chord Ring	5.63	121.1	21.5	112.6
Twenty-Chord Ring	3.34	65.7	19.7	66.8
Twelve-Chord Ring	4.13	–	–	82.6
Leader Cycle	2.31	40.9	17.7	46.2

Given the limited information available to human or agent participants in this problem, the decision process is relatively simple. Given the number of other linked participants with the same (conflicting) colour, and the number with a different (acceptable) colour, at any given point in time each participant makes a binary choice to "retain" or "switch" their existing colour. How did the human participants in the study by Kearns *et al.* (2006) make these decisions? It is difficult to be certain, since there is no published data on this. However, it seems likely that participants would repeatedly review the "badness" of the current situation, and have a probability of switching based on that badness. Clearly, conflicts are bad, but there may also be a tendency for the highly connected individuals in the "leader cycle" network to consider their choice as a leadership decision, thereby considering switching more negatively than other participants. We therefore assume a badness score B of the form:

$$B_{\text{retain}} = p + s \qquad (1)$$

$$B_{\text{switch}} = p + d + qn^r$$

where p, q, and r are numeric parameters, s is the number of linked participants with the same colour, d the number with a different colour, and $n = s + d$ the total number of links. It is therefore bad to retain the existing colour if there are many linked participants with the same colour, and it is bad to switch if there are many with a different colour, or if there are simply many links. The use of two parameters (q and r) to control the many-links effect means that very little is assumed about the form that it takes. Furthermore, Equation (1) incorporates all the information available to agents about the present state of affairs, although it does not include historical information about past actions.

The agents we have implemented make their "retain" or "switch" decisions probabilistically, choosing options based on the reciprocal of the badness score.

Tables 2, 3, and 4 show the simulation results when the p-q-r parameter space is explored (note that when $q = 0$, r becomes irrelevant; also, additional results for $p = 0.1$ are not shown – they are similar to those for $p = 0$). The time ratio for the two cases of the simple ring and the "leader cycle" is shown in Table 2. The human data in Table 1 suggests a ratio of 4.23 for these cases, and the only parameter combinations which approximate this are:

- $p = 0$, $q = 0.1$, $r = 2$;
- $p = 0$, $q = 1$, $r = 2$; and
- $p = 1$, $q = 1$, $r = 2$.

Table 3 shows a similar set of data for the time ratio of the twelve-chord ring and the "leader cycle." The human data in Table 1 suggests a ratio of 1.79 for these cases, and the only parameter combination which approximates this is $p = 0$, $q = 0.1$, and $r = 2$. Thus two of the parameter combinations suggested by Table 2 are ruled out. The remaining parameter combination, which best fits the available human data, corresponds to a quite substantial many-links effect. The factor qn^r in Equation (1) is only 0.9 for an agent with three links, but 36.1 for the highly connected "leaders" in the "leader cycle," which have 19 links each. The selected parameter combination suggests that these "leaders" are thus assigning a very substantial "badness" to switching colours – perhaps partly influenced by the very use of the term "leader" for these well-connected individuals.

The two rings are not directly compared: a 2.36 ratio is expected, and from Tables 2 and 3 this must be roughly achieved (for the given p, q, r, the actual ratio is 3.60).

Table 2: Simulation data for network-colouring agents, showing the ratio for time taken with the simple ring to time taken with the "leader cycle." A value of 4.23 is suggested by the human data in Table 1.

Simple ring/leader cycle time ratio							
		q					
	r	**0**	**0.01**	**0.03**	**0.1**	**0.3**	**1**
***p* = 0**	**0**	17.7	17.4	15.3	21.3	24.4	45.7
	0.1		15.0	16.2	21.2	24.6	45.9
	0.2		15.7	16.1	24.2	28.5	48.8
	0.25		18.3	18.3	23.1	29.6	44.7
	0.33		18.5	19.2	20.7	27.7	41.9
	0.5		16.0	19.1	20.5	28.4	38.6
	1		17.7	18.8	25.0	27.7	22.6
	2		15.9	9.6	**5.5**	32.9	**2.7**
***p* = 1**	**0**	17685	17991	18810	19962	15316	13757
	0.1		16833	17738	13738	17534	13070
	0.2		17002	15615	17431	22864	14508
	0.25		17659	20327	17225	16105	13260
	0.33		15411	16425	15406	19527	10508
	0.5		22049	17335	18417	20160	8786
	1		16949	15985	17826	14269	6986
	2		18435	17145	3507	696	**3.8**

Finally, Table 4 shows actual times taken (in simulation steps) for the "leader cycle." The two parameter combinations $p = 0$, $q = 1$, $r = 2$ and $p = 1$, $q = 1$, $r = 2$ suggested by Table 2 (but rejected by Table 3) are seen to be impossibly slow. The parameter combination $p = q = r = 0$, corresponding to discarding the many-links effect, is 22 times faster than the parameter combination $p = 0$, $q = 0.1$, $r = 2$ selected above.

Table 3: Simulation data for network-colouring agents, showing the ratio for time taken with the twelve-chord ring to time taken with the "leader cycle." A value of 1.79 is suggested by the human data in Table 1.

Twelve-chord ring/leader cycle time ratio							
		q					
	r	**0**	**0.01**	**0.03**	**0.1**	**0.3**	**1**
***p* = 0**	**0**	4.5	3.9	4.4	4.5	4.8	6.8
	0.1		4.6	4.8	4.7	4.9	6.6
	0.2		4.0	4.4	5.1	4.2	7.0
	0.25		5.1	4.7	4.7	5.6	6.6
	0.33		4.4	4.0	3.8	6.5	7.4
	0.5		4.3	5.2	4.8	5.1	7.0
	1		4.3	4.4	5.6	6.8	5.3
	2		4.6	2.6	**1.5**	1.0	1.1
***p* = 1**	**0**	77	43	36	33	25	21
	0.1		38	41	25	29	14
	0.2		41	33	31	30	20
	0.25		34	45	32	29	17
	0.33		35	28	29	30	18
	0.5		42	34	33	25	16
	1		302	37	24	16	13
	2		32	22	4.9	2.6	2.6

Table 4: Simulation data for network-colouring agents, showing the actual time taken (in simulation steps) for the "leader cycle." The parameter combination $p = q = r = 0$, corresponding to an absence of the many-links effect, is 22 times faster than the parameter combination $p = 0$, $q = 0.1$, $r = 2$.

Time in simulation steps for leader cycle							
		q					
	r	**0**	**0.01**	**0.03**	**0.1**	**0.3**	**1**
***p* = 0**	**0**	4.3	4.3	4.4	4.2	4.5	6.7
	0.1		4.3	4.3	4.7	4.7	7.7
	0.2		4.6	4.2	4.3	4.8	8.3
	0.25		4.4	4.3	4.6	5.0	8.9
	0.33		4.3	4.2	4.5	5.5	10.2
	0.5		4.4	4.1	4.8	6.1	15.1
	1		4.5	4.6	5.6	10.9	61
	2		6.0	17.1	**92**	968	231,798
***p* = 1**	**0**	38	36	36	35	34	26
	0.1		38	38	41	33	25
	0.2		39	41	36	28	23
	0.25		44	36	37	28	24
	0.33		40	41	38	27	26
	0.5		34	39	33	25	27
	1		37	38	29	25	63
	2		32	30	92	651	242,026

If our selected parameter combination corresponds to actual human decision-making, it seems that people are solving the problem with a less-than-optimal approach. This is supported by the fact that there are several network topologies (which we have not examined here) for which groups of humans solve the problem only about half the time (Kearns *et al*. 2006). It seems that the approximately linear relationship between average distance and time highlighted in Table 1 is the result of a human decision process which unnecessarily slows down networks containing highly connected people.

How much confidence can we have in our agent-based model of network colouring? A reasonable degree of confidence is justified in this case. The network-colouring task is fairly simple, and the restricted information available means that human decision-making for the problem must use the same information used in Equation (1), namely the number of links to people with the same colour, and the number to people with a different colour. The mathematical formula we used in Equation (1) is quite general in combining these two items of information, and exploring the p-q-r parameter space finds only one combination of parameters ($p = 0$, $q = 0.1$, $r = 2$) consistent with the empirical data. It therefore seems very likely that the many-links effect is real, even if it may not have precisely the same form as Equation (1). However, it would be desirable to have additional empirical data to confirm this conservatism of highly connected people, and to confirm that the omission of historical information from Equation (1) is not a serious limitation.

We used the simulation model shown in Figure 4 to explore the parameter space. For the alternative implementation shown in Figure 3, which allows interaction between software agents and a human participant, we substituted the $p = 0$, $q = 0.1$, $r = 2$ parameter combination into Equation (1), in order to make the agents as human-like as possible. This permits experiments in which a single human being plays the role of a highly connected individual, while software agents play the other roles. For additional realism, we also added a text communication facility, where agents could report on the reasons for their decision-making. This facility is shown at the bottom of Figure 3.

As noted in the Introduction, since the model outputs have been used to select parameter values, they cannot also be used as a validation technique (although, as we have seen, a reasonable degree of confidence is justified in the model). However, we are planning to use the software tool shown in Figure 3 (as well as future refinements of it) to collect additional empirical data for further model validation.

4. DISCUSSION

We have explored parameter selection in an agent-based model inspired by the laboratory experiments of Kearns *et al*. (2006) on network colouring. To calibrate this model, we chose model parameters indirectly, selecting that combination of parameters for which the model output best matched the empirical linear relationship between average distance and time. Such indirect parameter selection means that model outputs cannot also be used for validation purposes. However, this need not be a major concern if there is independent evidence supporting the model design, or if the design is sufficiently general that no unjustified assumptions are made.

Our indirect approach to parameter selection generated an unexpected insight. The parameters best fitting the empirical data corresponded to a moderately high degree of conservatism by highly connected individuals, and suggested that solution times were up to 22 times as long as they needed to be because of this.

Additional empirical data on the behaviour of highly-linked individuals would be desirable. However, a reasonable degree of confidence is justified in this insight, since model design decisions were either justified on the basis of empirical data, or kept very general, with the choice of parameters being decided (indirectly) by comparing model outputs to empirical data. Fortunately, the empirical data did not lead to any ambiguity about the choice of model parameters.

Because the indirect approach requires a single “maximum” in terms of degree of fit of model outputs to empirical data, it is best suited to models with relatively few parameters. Efficiency concerns, given the need to explore the entire parameter space, reinforce this limitation. With appropriate care, the indirect approach to parameter selection therefore seems to have considerable potential value in building social agent-based models, as long as they are relatively simple in design.

We intend to build on the work presented in this paper by conducting further human experiments with the network colouring task, revising the agent-based simulation in the light of those experiments, and applying the resulting insights to the larger project of developing realistic agent-based models which can interact with human beings in an experimental or game setting (Dekker 2008). Understanding the situations in which people behave differently from optimal computer algorithms will be particularly important in that activity.

REFERENCES

Cohen, R.J., Swerdlik, M.E., & Phillips, S.M. (1988) *Psychological Testing and Assessment: An Introduction to Tests and Measurement*, 3rd edition, Mountain View, California: Mayfield.

Crooks, A., Hudson-Smith, A., & Dearden, J. (2009) “Agent Street: An Environment for Exploring Agent-Based Models in Second Life,” *Journal of Artificial Societies and Social Simulation*, Vol. 12, No. 4, Paper 10, jasss.soc.surrey.ac.uk/12/4/10.html

Dekker, A.H. (2007) “Studying Organisational Topology with Simple Computational Models,” *Journal of Artificial Societies and Social Simulation*, Vol. 10, No. 4, Paper 6, jasss.soc.surrey.ac.uk/10/4/6.html

Dekker, A.H. (2008), “Experiments in Planning and the ‘ISLAND MISSION’ Game,” *SimTecT 2008 Simulation Technology and Training Conference*.

Dekker, A.H. (2009) “Deriving Agent-Based Simulation Models from Textual Data,” in Anderssen, R.S., Braddock, R.D., and Newham, L.T.H. (eds) *18th World IMACS Congress / MODSIM09 International Congress on Modelling and Simulation*, Modelling and Simulation Society of Australia and New Zealand, pp 2936–2941, www.mssanz.org.au/modsim09/H4/Dekker_H4.pdf

Epstein, J.M. (2008) “Why Model?,” *Journal of Artificial Societies and Social Simulation*, Vol. 11, No. 4, Paper 12, jasss.soc.surrey.ac.uk/11/4/12.html

Gibbons, A (1985) *Algorithmic Graph Theory*, Cambridge: Cambridge University Press.

Gilbert, N. (2008) *Agent-Based Models*, London: SAGE Publications.

Kearns, M., Suri, S., & Montfort, N. (2006) “An Experimental Study of the Coloring Problem on Human Subject Networks,” *Science*, Vol. 313, 11 August, pp 824–827.

Moss, S. (2008) “Alternative Approaches to the Empirical Validation of Agent-Based Models,” *Journal of Artificial Societies and Social Simulation*, Vol. 11, No. 1, Paper 5, jasss.soc.surrey.ac.uk/11/1/5.html

Özman, M. (2008) “Network Formation and Strategic Firm Behaviour to Explore and Exploit,” *Journal of Artificial Societies and Social Simulation*, Vol. 11, No. 1, Paper 7, jasss.soc.surrey.ac.uk/11/1/7.html

Windrum, P., Fagiolo, G., & Moneta, A. (2007) “Empirical Validation of Agent-Based Models: Alternatives and Prospects,” *Journal of Artificial Societies and Social Simulation*, Vol. 10, No. 2, Paper 8, jasss.soc.surrey.ac.uk/10/2/8.html